

A close-up, high-angle photograph of a speaker driver, likely a tweeter or midrange driver, showing the internal cone and suspension. The speaker is dark-colored with a metallic-looking center. The background is black.

The Complete Guide to High-End Audio

Fifth Edition

Robert Harley

Editor-in-Chief
The Absolute Sound

The Complete Guide to High-End Audio

Fifth Edition

Robert Harley

Acapella Publishing

Carlsbad, California

hifibooks.com

Copyright © 1994—2015 by Robert Harley

Fifth Edition

First Printing—2015

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system without written permission from the publisher, except for the inclusion of brief quotations in a review.

International Standard Book Number:

978-0-9786493-6-4

Cover Concept and Design: *Torquil Dewar*

Printed in the United States of America

Trademark Notice

This book comments on and describes various audio and video components. Many such products are identified by their tradenames. In most cases, these designations are claimed as legally protected trademarks by the companies manufacturing the products. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark. You should investigate a claimed trademark before using it for any purpose other than to refer to the product.

Warning Disclaimer

This book is designed to provide the most current, correct, and clearly expressed information regarding the subject matter covered. It is sold with the understanding that the publisher and author are not engaged in rendering audio/video system design.

It is not the purpose of this guide to reprint all the information otherwise available to the author and/or publisher, but to complement, amplify, and supplement other texts. You are urged to read all available material and learn as much as possible about audio, video, and home theater to tailor the information to your individual needs. For more information, readers are encouraged to consult with professional specialty audio and video retailers or manufacturers for information concerning specific matters before making any decision.

Every effort has been made to make this guide as complete and accurate as possible. However, there may be mistakes both typographical and in content. Therefore, this text should be used as a general guide and not as the ultimate source of audio and video information. Furthermore, this text contains information current only up to the printing date.

The purpose of this guide is to educate and entertain. The author and Acapella Publishing shall have neither liability nor responsibility to any person or entity with respect to any loss or damage caused, or alleged to be caused, directly or indirectly by the information contained in this book.

If you do not wish to be bound by the above, you may return this book to the place of purchase for a full refund.

Contents

Foreword

Preface to the Fifth Edition

About the Author

1 What Is High-End Audio?

2 Choosing a High-End Audio System

Introduction

Choosing the System Best Suited to Your Needs

Setting Your Budget

The Complete vs. the Incremental Purchase

 Value vs. Luxury Components

Allocating Your Budget to Specific Components

Upgrading a Single Component

How to Read Magazine Reviews

System Matching

Do's and Don'ts of Selecting Components

Your Relationship with the Retailer

Used Equipment

Product Upgrades

Component Selection Summary

3 Becoming a Better Listener

Audiophile Values

Pitfalls of Becoming a Critical Listener

Sonic Descriptions and their Meanings

 Tonal Balance

 Overall Perspective

 The Treble

 The Midrange

 The Bass

 Soundstaging

Dynamics

Detail

Pace, Rhythm, and Timing

Coherence

Musicality

Notes on Learning Descriptive Terms

Critical-Listening Setup Procedures

Single-Presentation Listening—What It's All About

Critical Listening Summary

Addendum: Level Matching

4 Preamplifiers

Introduction

How to Choose a Preamplifier

Balanced and Unbalanced Connections

What to Listen For

Tubes vs. Transistors

Tube Life and Replacement Options

The Line-Stage Preamplifier

The Phono-Stage Preamplifier

RIAA Equalization

Phono-Stage Gain

Cartridge Loading

The Digital Preamplifier

Audio/Video Controllers and Multichannel Preamplifiers

How a Preamplifier Works

Balanced and Unbalanced Preamplifiers

5 Power and Integrated Amplifiers

Introduction

How to Choose a Power Amplifier

Integrated Amplifiers

How Much Power Do You Need?

The dBW Power Rating

Why Amplifier Output Current Matters

What to Look For when Comparing Power Ratings

- Why Amplifier Power Isn't Everything
- Other Power-Amplifier Considerations
 - Tubes vs. Transistors
 - Balanced Inputs
 - Bridging
 - Bi-Amping
- What to Listen For
- A Survey of Amplifier Types
 - Single-Ended-Triode Amplifiers
 - Single-Ended Solid-State Amplifiers
 - Class D (Switching) Power Amplifiers
 - Digital Amplifiers
 - A Unique Amplifier
- Output-Stage Topology and Class of Operation
- How a Power Amplifier Works
 - The Power Supply
 - Input and Driver Stages
 - Output Stage
 - How a Tubed Power Amplifier Works

6 Loudspeakers

- Introduction
- How to Choose a Loudspeaker
 - Other Guidelines in Choosing Loudspeakers
 - Finding the Right Loudspeaker—Before You Buy
- What to Listen For
- Loudspeaker Types and How They Work
 - The Dynamic Driver
 - The Coaxial (Concentric) Driver
 - Dynamic Compression
 - Problems with Dynamic Drivers
 - The Electromagnetic Dynamic Driver
 - The Planar-Magnetic Transducer
 - The Heil Air-Motion Transformer
 - The Electrostatic Driver

The Dipolar Radiation Patterns of Ribbons and Electrostatics

Bipolar and Omnidirectional Loudspeakers

The Horn Driver

Waveguides

The Plasma Driver

The Full-Range Driver

Loudspeaker Enclosures

Infinite-Baffle Loading

Reflex Loading

Passive Radiators

Transmission-Line Loading

Isobarik Loading

The Finite Baffle

System Q

Powered and Servo-Driven Woofers

Enclosure Resonances

Enclosure Shapes

Crossovers

Digital Loudspeakers

DSP Loudspeakers

Subwoofers

Subwoofer Technical Overview

Loudspeaker Stands

7 Digital Source Components, Part One: Disc Players, DACs, and Interfaces

Introduction

How to Choose a Digital Source—Overview

What to Listen For

Disc Players

Super Audio Compact Disc (SACD)

Universal Disc Players

Blu-ray Disc

Bass Management in SACD and Universal Players

How to Choose a Digital Source—Features and Specs

Disc Player and DAC Features and Specs

Digital Interfaces

The S/PDIF Interface

The USB Interface

Ethernet

The I²S Interface

DSD over PCM (DoP)

Proprietary Digital Interfaces

Wireless Connection

Jitter in the Digital Interface

Outboard Clocks

Asynchronous Sample-Rate Conversion

High-Resolution Digital Audio: Why 44.1kHz Sampling and 16-bit Quantization Aren't Enough

16 Bits, 20 Bits, 24 Bits

A Caveat About High-Resolution Digital Audio

Technical Addenda

Direct Stream Digital (DSD) Encoding on SACD

High Definition Compatible Digital (HDCD)

How a DAC Works

Custom Digital Filters

Digital-to-Analog Conversion

Delta-Sigma DACs

Analog Stages

Balanced Outputs

8 Digital Source Components Part Two: Music Servers

Music Server Sound Quality: Better than Disc

Downloads

Download Caveats

Turnkey Music Servers vs. Computer-Based Servers

The Turnkey Music Server

Another Type of Turnkey Music Server

Turnkey Server Considerations

Using a Smartphone or Tablet with a Music Server

Internet Radio and Music Subscription Services

Metadata

File Formats

The Do-It-Yourself, PC-Based Music Server

The Importance of Bit Transparency

Computer Requirements

Operating System and Playback Software

Hard-Disk Storage

Loading Your Server with Music

Getting Music from the Server to Your Playback System

Networking

A Simple Do-It-Yourself Music Server

Transferring an LP Collection to a Server

9 Turntables, Tonearms, and Cartridges: The LP Playback System

Introduction

LP Playback Hardware Overview

System Hierarchy: Why the LP Front End Is So Important

How to Choose an LP Playback System

What to Listen For

Technical Aspects of Choosing an LP Front End

The Turntable

The Base and Plinth

Sprung and Unsprung Turntables

The Platter and Bearing Assembly

Platter Mats, Record Clams, and Vacuum Hold-Down Systems

The Drive System

The Tonearm

The Phono Cartridge

Moving-Magnet and Moving-Coil Cartridges

The Strain-Gauge Cartridge

Stylus Shapes and Cantilever Materials

LP Playback System Setup

Record Care and Cleaning

Vinyl as Art: Half-Speed Mastering, 45rpm Pressings, 180-gram Vinyl, and Direct-to-Disc LPs

10 Cables and Interconnects

Introduction

How to Choose Cables and Interconnects

How Much Should You Spend on Cables and Interconnects?

What to Listen For

Binding Posts and Cable Terminations

Bi-Wired Loudspeaker Cables

Balanced and Unbalanced Lines

Cable and Interconnect Construction

Conductors

The Dielectric

Terminations

Geometry

Terminated Cables and Interconnects

Battery Bias in Cables and Interconnects

Cable and Interconnect Specifications

Cables in the Power Amplifier/Loudspeaker Interface

11 Audio for Home Theater

Introduction

Overview of Home-Theater Systems

Should You Choose a 5.1-, 7.1-, or 9.1-Channel System?

Home-Theater Controllers

How to Choose a Controller

Inputs, Outputs, and Source Switching

Automatic Calibration

DSP Speaker and Room Correction

Analog Bypass Modes

Bass-Management Flexibility

8-Channel Analog Input

Digital-to-Analog Conversion

7.1-Channel and 9.1-Channel Playback from 5.1-Channel Sources

THX Certification

3D Capability, Network Connection, Wireless Streaming, and Multi-Zone
Multichannel Power Amplifiers

How to Choose a Home-Theater Power Amplifier

Loudspeakers for Home Theater

The Center-Channel Speaker

Adding a Center Speaker to Your System

Left and Right Speakers

Surround Speakers

Dipolar and Bipolar Surround Speakers

Surround-Back Speakers

Height Speakers

Subwoofers

Setting Up a Home Theater

Basic Setup

Acoustical Treatment

Speaker Placement

Surround Speaker Placement

Center Speaker Placement

Left and Right Speaker Placement

Calibrating a Home Theater

Bass Management

Setting Individual Channel Levels

Adding Home Theater without Compromising Music Performance

Addendum: Surround-Sound Formats Explained

12 Multichannel Audio

Introduction

A Short History of Multichannel Audio

Do We Want Multichannel Music Playback?

How to Get Multichannel Audio in Your Home

Loudspeaker Types and Placement

Bass Management

Calibration

Multichannel Playback from 2-Channel Sources

Ambisonics

13 Personal Audio: Headphones, Desktop, and Mobile Audio

Introduction

Headphones

Earbuds, Earphones, and In-Ear Monitors

Headphones

Dynamic, Electrostatic, and Planar-Magnetic Headphones

Headphone Enhancement Technologies

Crossfeed Circuits

Smyth Virtual Surround (SVS)

Binaural Listening

Multichannel Headphones

Headphone Specifications

Headphone Amplifiers

Balanced and Unbalanced Drive

Portable Audio

Desktop Audio

14 System Setup Secrets, Part One: Loudspeaker Placement and Room Acoustics

Introduction

Loudspeaker Placement

Loudspeaker Placement in Asymmetrical Rooms

Short-Wall vs. Long-Wall Placement

Dipolar and Bipolar Loudspeaker Placement

Subwoofer Setup and Placement

Multichannel Loudspeaker Placement

Loudspeaker Placement Summary

Common Room Problems and How to Treat Them

Acoustical Dos and Don'ts

A Short Course in Acoustical Theory

Listening-Room Resonance Modes

Optimizing Dimensional Ratios

Standing Waves

Reverberation

Building a Listening Room from Scratch

Isolating the Listening Room

The Walldamp Technique

DSP Room Correction

15 System Setup Secrets, Part Two: Expert Tuning Techniques, Accessories, Equipment Racks, and AC Power Conditioners

Introduction

Accessories

How to Choose Accessories

Analog Accessories

System Setup Techniques

Equipment Placement

Cables and Interconnects

Tubed Equipment

LP Playback

Digital Playback

Loudspeakers

AC Power

General

Equipment Racks and Isolation Devices

Equipment Racks

Vibration-Control Accessories: Spikes, Feet, and Cones

AC Power Conditioners and AC Cords

AC Power Cords

Dedicated AC Lines

A Final Note

16 Specifications and Measurements

Preamplifiers

Power Amplifiers

Loudspeakers

Digital Components

Appendix A: Sound and Hearing

Introduction: What is Sound?

Period and Frequency

Wavelength

- Phase
- Absolute Polarity
- Complex Waves
- Comb Filtering
- Absorption, Reflection, and Diffusion
- Diffraction
- The Decibel (dB)
- Frequency, Loudness, and Equal Loudness Contours
 - Weighting Filters
- The Dynamic Nature of Music
- Localization
- Other Psychoacoustic Phenomena

Appendix B: Audio and Electronics Basics

- Voltage, Current, Resistance, and Power
 - Ohm's Law
 - Series and Parallel Circuits
- Alternating Current (AC)
- Electromagnetic Induction, Inductance, and Capacitance
 - Filters
- Impedance
 - Loudspeaker Phase Angle
 - Capacitor Types
- Anatomy of an Audio Product
 - The Power Supply
 - Amplifier Circuits
 - Amplifier Distortion
 - Vacuum Tube Amplifiers
 - Operational Amplifiers
- Digital Electronics

Appendix C: Digital Audio Basics

- Introduction
- Binary Number System
- Sampling and Quantization
 - Sampling Rate, Nyquist Theorem, and Aliasing

Quantization

Dither

Digital Audio Storage

Error Correction

Digital-to-Analog Conversion

Jitter Explained

Digital Signal Processing (DSP)

Perceptual Coding

Glossary

Index

Foreword

By Keith Jarrett

Music is the sonic motion of intention. With words, sound can be divorced from meaning by taking away the physical quality of speech. But music's meaning is *in* its physical quality: its sound. When a musician plays something a certain way and we can't hear the intent (the reason) behind it, we are hearing wasted motion, and register it as such because we haven't been given enough clues about the intent. We can then grow to think that everything is only gesture, and miss the real thing.

The media through which we hear music (our systems, rooms, etc.) cannot be separated from our ability to experience the music. It isn't the same *music* on a different system because we cannot separate music's rhetoric (its words) from its physical reality (its delivery). This makes the "delivery systems" (our stereos) more important than we might think they are. Can they *tell* us what the musicians on the recording are telling us?

As a musician, I often—too often—had the following experience: I would play a concert, hear the tape afterward, and wonder what was missing. I would remember incredible things in the concert that just weren't on the tape. The notes were there, but notes are not music. Where was the music, the *intention*?

We could think of it this way: On the tape, the rhetoric had no meaning. Had I trusted the tape and not my memory of the actual event, I would have never grown to understand that, even though the sound is on tape, it doesn't mean you've recorded the *music*. If you've heard a certain CD on a certain system, it doesn't necessarily mean that you've heard what's on the CD. We must learn to trust the responses of our own system—our ears—to music systems. Of course, this demands that we be in touch with ourselves—no easy thing.

People to whom music is important need to get close to the intention in a recording, and there's only one way to do this in the home: learn about the world of audio equipment. Use your (and others') ears to help remove whatever hinders you from the musical experience on the recording. Of course, it's not only the

reproduction side that needs care—but that's the only side the listener has control over.

For instance, it's demonstrable that by merely flipping a two-pronged AC plug on a CD player, or even a turntable, a record you thought you didn't like can become a favorite—just because the polarity was wrong. Since music cannot be divorced from its emotional content, the *sound* of a record can determine whether you think you like the *music*. And vice versa, when you can't listen to music you really think you like because of how it was *recorded*.

Obviously, the musical experience is a delicate, complex thing, and we humans are more sensitive than we sometimes think. But we have the option to tune our music systems to better balance the equation. We can get closer to what we want if we *know* what we want.

There are stereo components that approximate the musical experience at many different price levels. We all know what our financial limitations are; but, given the desire to improve our systems, we *can* do it.

It by no means follows that musicians have to be audiophiles. Though I've been recording since 1965, I didn't seriously think about much of this until the mid-1980's. But audiophiles and music lovers push the envelope, and we all benefit. Also, the more serious audiophiles are determined to keep their minds and ears open, keep learning, and try to remain patient during the process. Doing this thing right can take time.

There are a lot of people out there listening to all of these components for us. I recommend *using* this fact, and *carefully* reading others' evaluations, until you can tell whether a reviewer's preferences in sound match your own priorities. You *can* sort of get to know these guys over a period of time.

But, of course, it's *your* ears that count. I think you should pay attention to their needs. After all, we're talking nutrition in an age of diet soft drinks.

Preface to the Fifth Edition

Each time I revise and update this book, I'm made vividly aware of the technological changes that shape the ways in which we enjoy music. Although the fourth edition was published less than five years ago, that short span of time has seen more profound changes in audio technology—transformative developments that open startling new possibilities for accessing music—than in any previous *decade*.

The new technologies that most readily come to mind are music servers, streaming audio, wireless networking, and high-resolution downloads. As I write this, I'm listening to full CD-quality streaming audio on my desktop system, and have the ability to instantly access just about any piece of music I can think of. It's like living in a virtually infinite music store. I can stream that music wirelessly to audio components distributed throughout my house. In my main audio system, a network-attached storage drive holds a growing collection of high-resolution downloads along with most of my CD collection, all accessible with a few finger taps on a tablet. That system includes an audio component that was unimaginable five years ago: an integrated amplifier and high-res wireless streaming DAC that's more accurately viewed as a general-purpose, software-controlled audio platform that can be programmed, via Internet download, into whatever configuration suits my needs. And when I travel, I can take high-res music with me, and listen to it through planar-magnetic headphones whose sound quality rivals that of six-figure loudspeakers. Never in history has so much music been so readily available to so many listeners.

But beyond these attention-grabbing examples of advancing technology, the quest for more lifelike musical reproduction has moved forward in myriad small ways across nearly every aspect of audio reproduction. These advances may not be obvious or profound, but, cumulatively, they take us that much closer to high-end audio's *raison d'être*: equipment that makes it possible to fully experience at home the intents of the composer and musicians. Take, for example, the apparently prosaic power amplifier: How much room for improvement was there in this category of audio component? As it turns out, a lot. Today's best power

amps have significantly advanced the state of the art with improved power supplies and novel circuit topologies.

Every advance in one product category spurs improvements in other areas—a more resolving amplifier provides designers of loudspeakers and source components with a more transparent window on the sounds of their own products. Even that most venerable of audio technologies, the LP record, has a few new tricks up its sleeve. With the market for vinyl playback gear and new LPs exploding, the industry has responded with a slew of innovative products aimed at extracting the most music from vinyl. Even more exciting, most of the records being pressed today sound better than LPs ever have. Audiophile reissues of classic albums, produced with loving devotion, bring the music back to life and engender the feeling that, for the first time, you're hearing the music as it was meant to be heard.

Audiophiles talk of the “golden age” of audio as being the 1950s or 1960s, when many of today's audio paradigms were established. As great as that period was, I think we're in the middle of a new golden age, and not just because of the advancing state of the art. The past few years have been marked by a democratization of high-end audio in which simple and affordable products designed and built with high-end values have become commonplace. Most of the people who buy these products wouldn't identify themselves as audiophiles, but these listeners nonetheless want high-quality sound. The future of high-performance audio looks bright indeed.

Finally, a few notes about how to use this book. As in the previous editions, each chapter begins with the most basic information and progresses to more technical discussions. When you've reached a level of complexity you're comfortable with and don't wish to go deeper, simply skip to the beginning of the next chapter, to begin reading the most basic information about the next subject. This book is meant to be used as a reference, rather than read straight through from beginning to end.

Keep in mind that you don't need to understand the technical aspects of audio to enjoy music in your home. I've included the more technical content for those readers interested in knowing how audio works, and to make this book truly *The Complete Guide to High-End Audio*.

Robert Harley
Carlsbad, California

About the Author

Robert Harley is the Editor-in-Chief of *The Absolute Sound* magazine. Founded in 1973, *The Absolute Sound* is the world's most respected journal of high-end audio. He is the author of two other books on home audio, *Home Theater for Everyone* and *Introductory Guide to High-Performance Audio Systems*. His more than 1000 published equipment reviews and articles on music and home-theater sound reproduction have helped thousands of enthusiasts improve their home-entertainment systems. His books have sold more than a quarter of a million copies in five languages.

Robert is also the editor of *The Absolute Sound's Illustrated History of High-End Audio*, a series of richly illustrated large-format books celebrating the audio industry's legendary designers and their iconic products.

Robert Harley holds a degree in recording engineering and has taught a college degree program in that field. He has worked as a recording engineer and studio owner, compact-disc mastering engineer, technical writer, and audio journalist. Before joining *The Absolute Sound* and *The Perfect Vision* in 1999, he was Technical Editor of *Stereophile* magazine for eight years, and also served in that capacity at *Fi: The Magazine of Music and Sound*.

1 What Is High-End Audio?

High-end audio is about passion—passion for music, and for how well it is reproduced. High-end audio is the quest to re-create in the listener’s home the musical message of the composer or performer with the maximum realism, emotion, and intensity. Because music is important, re-creating it with the highest possible fidelity is important.

High-end audio products constitute a unique subset of music-reproduction components that bear little similarity to the “stereo systems” sold in department stores. A music-reproduction system isn’t a home appliance like a washing machine or toaster; it is a vehicle for expressing the vast emotional and intellectual potential of the music encoded on our records and CDs. The higher the quality of reproduction, the deeper our connection with the music.

The high-end ethos—that music and the quality of its reproduction matter deeply—is manifested in high-end audio products. They are designed by dedicated enthusiasts who combine technical skill and musical sensitivity in their crafting of components that take us one step closer to the original musical event. High-end products are designed by ear, built by hand, and exist for one reason: to enhance the experience of music listening.

A common misperception among the hi-fi-consuming public is that high-end audio means *high-priced* audio. In the mass-market mind, high-end audio is nothing more than elaborate stereo equipment with fancy features and price tags aimed at millionaires. Sure, the performance may be a little better than the hi-fi you find at your local appliance store, but who can afford it? Moreover, high-end audio is seen as being only for trained, discriminating listeners, snobs, or gadget freaks—not for the average person on the street.

High-end audio is none of these things.

First, the term *high-end* refers to the products’ *performance*, not their price. Many true high-end systems cost no more—and often less—than the all-in-one rack systems sold in department stores. I’ve heard many inexpensive systems that

capture the essence of what high-quality music reproduction is all about—systems easily within the budgets of average consumers. Although many high-end components *are* high-priced, this doesn't mean that you have to take out a second mortgage to have high-quality music reproduction in your home. A great-sounding system can be less expensive than you might think.

Second, high-end audio is about communicating the musical experience, not adding elaborate, difficult-to-operate features. In fact, high-end systems are much easier to use than mass-market mid-fi systems. This is because the high-end ethic eliminates useless features, instead putting the money into sound quality. High-end audio is for music lovers, not electronics whizzes.

Third, *anyone* who likes music can immediately appreciate the value of high-quality sound reproduction. It doesn't take a "golden ear" to know what sounds good. The differences between good and mediocre music reproduction are instantly obvious. The reaction—usually pleasure and surprise—of someone hearing a true high-end audio system for the first time underscores that high-end audio can be appreciated by everyone. If you enjoy music, you'll enjoy it more through a high-end system. It's that simple.

Finally, the goal of high-end audio is to make the equipment "disappear"; when that happens, we know that we have reached the highest state of communication between musician and listener. High-end audio isn't about equipment; it's about music.

The high-end credo holds that the less the musical signal is processed, the better. Any electronic circuit, wire, tone control, or switch degrades the signal—and thus the musical experience. This is why you won't find graphic equalizers, "spatial enhancers," "subharmonic synthesizers," or other such gimmicks in high-end equipment. These devices are not only departures from musical reality, they add unnecessary circuitry to the signal path. By minimizing the amount of electronics between you and the musicians, high-end audio products can maximize the directness of the musical experience. Less is more.

Imagine yourself standing at the edge of the Grand Canyon, feeling overwhelmed by its grandeur. You experience not only the vastness of this massive sculpture carved deep into the earth, but all its smaller features jump out at you as well, vivid and alive. You can discern fine gradations of hue in the rock layers—distinctions between the many shades of red are readily apparent. Fine details of the huge formations are easily resolved simply by your looking at them, thus deepening your appreciation. The contrasts of light and shadow highlight the apparently infinite maze of cracks and crevasses. The longer and closer you

look, the more you see. The wealth of sensory input keeps you standing silently at the edge, in awe of nature's unfathomable beauty.

Now imagine yourself looking at the Grand Canyon through a window made of many thicknesses of glass, each one less than perfectly transparent. One pane has a slight grayish opacity that dulls the vivid hues and obliterates the subtle distinctions between similar shades of color. The fine granular structure of the next pane diminishes your ability to resolve features in the rock. Another pane reduces the contrast between light and shadow, turning the Canyon's immense depth and breadth into a flat canvas. Finally, the windowframe itself constricts your view, destroying the Canyon's overall impact. Instead of the direct and immediate reality of standing at the edge of the Grand Canyon, what you see is gray, murky, lifeless, and synthetic. You may as well be watching it on television.

Hearing reproduced music through a mediocre playback system is like looking at the Grand Canyon through those panes of glass. Each component in the playback chain—CD player, turntable, preamplifier, power amplifier, loudspeakers, and the cables that connect them—in some way distorts the signal passing through it. One product may add a coarse, grainy character to instrumental textures. Another may reduce the dynamic contrasts between loud and soft, muting the composer's or performer's expression. Yet another may cast a thick, murky pall over the music, destroying its subtle tonal colors and overlaying all instruments with an undifferentiated timbre. Finally, the windowframe—that is, the electronic and mechanical playback system—diminishes the expanse that is the musicians' artistic intent.




Fig.1-1 Each component in an audio system can be thought of as a piece of glass through which we experience music. (Courtesy AudioQuest)

High-end audio is about removing as many panes of glass as possible, and making those that remain as transparent as they can be. The fewer the panes, and the less effect each has on the information passing through it, the closer we get to the live experience and the deeper our connection with the musical message.

Why are high-end audio products more transparent windows on the musical event than mass-market “stereo systems”? High-end products are designed to *sound* good—that is, like the real thing. They’re not necessarily designed to perform “well” according to some arbitrary technical specification. The true high-end designer *listens* to the product during its development, changing parts and trying different techniques to produce the most realistic sound possible. He combines technical skill with musical sensitivity to create a product that best conveys the musical experience. This dedication often becomes a zealous pursuit, involving many hundreds of listening hours and painstaking attention to every factor that influences the sound. Often, a more expensive part will be included to improve the product’s sound, while the retail price remains the same. The higher cost of this musically superior part comes off the company’s bottom line. Why? Because the high-end designer cares deeply about music and its reproduction.

Conversely, mass-market audio components are often designed to look good “on paper”—on the specification sheet—sometimes at the expense of actual sound quality. A good example of this is the “THD wars” of the 1970s and ’80s. THD stands for Total Harmonic Distortion, a specification widely used by uneducated consumers as a measure of amplifier quality. (If you’ve done this, don’t worry; before I learned more about audio, I, too, looked at THD figures.) The lower the THD, the better the amplifier was perceived to be. This led the electronics giants to produce products with vanishingly low THD numbers. It became a contest to see which brand had the most zeros after the decimal point in its THD specification (0.001%, for example). Many buyers bought receivers or amplifiers solely on the basis of this specification.

Although low THD is a worthy design goal, the problem arose in *how* those extremely low distortion figures were obtained. A technique to reduce distortion in amplifiers is called *feedback*—*taking* part of the output signal and feeding it back to the input. Large amounts of feedback reduce THD, but cause all kinds of other problems that degrade the amplifier’s musical qualities. Did the electronics giants care that the large amounts of negative feedback induced to reduce their products’ THD measurements actually made those products sound *worse*? Not a chance. The only thing that mattered was making a commodity that would sell in greater quantity. They traded musical performance for an insignificant technical specification that was sold to the public as being important. Those buyers choosing components on the basis of a specification sheet rather than listening ended up with poor-sounding systems. Ironically, the amplifiers that had the lowest THDs probably had the lowest quality of sound as well.

This example illustrates the vast difference between mass-market manufacturers’ and high-end companies’ conceptions of what an audio component should do. High-end manufacturers care more about how the product sounds than about how it performs on the test bench. They know that their audience of musically sensitive listeners will buy on the basis of sound quality, not specifications.

High-end products are not only designed by ear, but are often hand-built by skilled craftspeople who take pride in their work. The assemblers are often audiophiles themselves, building the products with as much care as if the products were to be installed in their own homes. This meticulous attention to detail results in a better quality of construction, or *build quality*. Better build quality can not only improve a product’s sound, but increase its long-term reliability as well. Moreover, beautifully hand-crafted components can inspire a

pride in ownership that the makers of mass-produced products can't hope to match.

High-end audio products are often backed by better customer service than mid-fi products. Because high-end manufacturers care more about their products and customers, they generally offer longer warranties, more liberal exchange policies, and better service. It is not uncommon for high-end manufacturers to repair products out of warranty at no charge. This isn't to say you should expect such treatment, only that it sometimes happens with high-end and is unthinkable with mass-market products. High-end companies care about their customers.

These attributes also apply to high-end specialty retailers. The high-end dealer shares a passion for quality music reproduction and commitment to customer service. If you're used to buying audio components at a mass-market dealer, you'll be pleasantly surprised by a visit to a high-end store. Rather than trying to get you to buy something that may not be right for you, the responsible high-end dealer will strive to assemble a system that will provide the greatest long-term musical pleasure. Such a dealer will put your musical satisfaction ahead of this month's bottom line.

Finally, most high-end products are designed and built in America by American companies. In fact, American-made audio components are highly regarded throughout the world. More than 40% of all American high-end audio production is exported to foreign countries, particularly the Far East. This is true even though high-end products cost about twice as much abroad as they do in the U.S., owing to shipping, import duties, and importer profit. The enthusiasm for American high-end products abroad is even more remarkable when one remembers the popular American misperception that the best audio equipment is made in Japan.

On a deeper level, high-end products are fundamentally different from mass-market products, in their conception, purpose, design, construction, and marketing. In all these differences, what distinguishes a high-end from a mass-market product is the designer's caring attitude toward music. He isn't creating boxes to be sold like any other commodity; he's making musical instruments whose performance will affect how his customers experience music. The high-end component is a physical manifestation of a deeply felt concern about how well music is reproduced, and, by extension, how much it is enjoyed by the listener.

The high-end designer builds products he would want to listen to himself. Because he cares about music, it matters to him how an unknown listener,

perhaps thousands of miles away, experiences the joy of music. The greater the listener's involvement in the music, the better the designer has done his job.

A digital-processor designer I know epitomizes this dedication. He had specified a premium-quality resistor at a certain point in his new design. This resistor cost \$1 rather than the pennies most resistors cost. Just as the design was about to go into production, he looked even harder for any changes that would improve the product's sound. For fun, he tried an exotic \$10 resistor in the circuit in place of the \$1 resistor. He was surprised at how much better the product sounded with this change, and couldn't bear to see the product shipped with the \$1 resistors. The company made the product with the \$10 resistors even though the retail price had already been established based on the parts cost using \$1 resistors. High-end designers try to add quality to, rather than subtract cost from, their products.

To the high-end designer, electronic or mechanical design isn't merely a technical undertaking—it's an act of love and devotion. Each aspect of a product's design, technical as well as musical, is examined in a way that would surprise those unaccustomed to such commitment. The ethos of music reproduction goes to the very core of the high-end designer's being; it's not a job he merely shows up for every day. The result is a much more powerful and intimate involvement in the music for the listener than is possible with products designed without this dedication.

What *is* high-end audio? What is high-end sound? It is when the playback system is forgotten, seemingly replaced by the performers in your listening room. It is when you feel the composer or performer speaking across time and space *to you*. It is feeling a physical rush during a musical climax. It is the ineffable roller-coaster ride of emotion the composer somehow managed to encode in a combination of sounds. It is when the physical world disappears, leaving only your consciousness and the music.

That is high-end audio.

2 Choosing a High-End Audio System

Introduction

Choosing a high-quality music-reproduction system is one of the most important purchasing decisions you'll make. Unlike buying home appliances, your selections in components will influence how deeply you appreciate and enjoy an art form—music. A great-sounding system can even change your lifestyle as music assumes a greater importance in your life. A hi-fi system is a vehicle for exploring the world of music; the better the system, the further and wider that vehicle will take you.

Although selecting hi-fi components may seem a daunting task, a little knowledge and preparation will go a long way toward realizing your dream system—and staying within your budget. The informed shopper knows that choosing the right components, matching those components to each other, and setting them up carefully are more important than a big bank account. This chapter will teach you to become a wise shopper and show you the path to assembling the most musically and aesthetically satisfying system possible for your money.

A high-end audio system is made up of many components, each with a specific job. These components generally fall into three basic categories: source components, control components, and playback components.

A source component is any device that recovers the audio signal from an information-storage medium or decodes that signal. Digital-to-analog converters (DACs), disc players, and turntables are all source components. Source components are also called the *front end* because they are the very beginning of the playback chain.

The control component is the preamplifier; it receives signals from the source components and controls which one is sent to the power amplifier for listening.

The preamplifier is the heart of a hi-fi system; all of the source components run through the preamplifier, which also controls volume and signal routing.

Playback components are the power amplifier and loudspeakers, which work together to convert the preamplifier's output into sound.

A traditional high-end audio system at one time consisted of a turntable, disc player, preamplifier with an integral phono stage, power amplifier, loudspeakers, and the cables to connect the components. Although this classic system architecture is still common, relatively new component categories offer many more choices in tailoring a system for your specific needs.

One of these component categories is the integrated amplifier with a built-in digital-to-analog converter (DAC) and possibly networking capability. The integrated amplifier accepts analog signals from analog source components, and digital datastreams from sources such as a computer's USB output or wirelessly from a mobile device or storage device connected to the network. Such a system could also include the ability to connect to your home's network and access music stored on other devices on that network. The integrated amplifier selects which source you want to hear, controls the volume, and drives your loudspeakers.

If you have no analog sources (e.g., your only source is a computer) you can choose a system with no preamplifier—a DAC with a built-in volume control drives a power amplifier directly. You don't need the integrated amplifier's analog inputs, ability to select between sources, and volume control. This approach saves money and results in a simpler system. (We'll look at these component categories in more detail in subsequent chapters.)

If you're a newcomer to audio, you may first want to read the introductory sections of [Chapters 4–13](#) before shopping for a system. Each of these chapters is devoted to a specific component, with an introductory section explaining the component's function, terminology, and features. Later sections in each chapter go into specific detail about selecting that component on musical and technical bases.

In this chapter, we'll look at how to choose a high-end system. The major topics include:

- *Choosing the system best suited to your needs*
- *Setting your budget*
- *Complete vs. incremental purchases*
- *Allocating your budget to specific components*

- *Upgrading a single component*
- *How to read magazine reviews of audio equipment*
- *System matching*
- *Dos and don'ts of selecting components*
- *Your relationship with the retailer*

Choosing the System Best Suited to Your Needs

Your first major decision will be whether your audio system will be used for 2-channel music listening, multichannel music, home theater, or a combination of all three. This fundamental choice will affect the types of components you choose as well as the allocation of your budget to specific components. We'll cover multichannel music and home-theater audio in their own chapters later, but keep in mind that spreading out your budget over five or seven loudspeakers and channels of amplification to drive them will necessitate buying products of lower quality than if you had invested your entire budget in just two loudspeakers and amplifier channels. Nonetheless, the general advice in this chapter on choosing an audio system applies equally well to any system configuration. Specific advice on choosing home-theater and multichannel audio equipment appears in [Chapters 11 and 12](#), respectively.

Just as a pickup truck is better suited to the farmer and a compact car to the city dweller, a hi-fi system ideal for a small New York City apartment would be entirely inadequate in a large suburban home. The hi-fi system must not only match your musical taste, as described in the next chapter, but must also suit your room and listening needs. (The following section is only an overview of how to choose the best system. More detailed information on how to select specific components is contained in [Chapters 4–13](#).)

Many of the guidelines are fairly obvious. First, match the loudspeaker size to your listening room. Large, full-range loudspeakers don't work well in small rooms. Not only are large loudspeakers physically dominating, they tend to overload the room with bass energy. A loudspeaker that sounds fine in a 17' by 25' room will likely be thick, boomy, and bottom-heavy in a 12' by 15' room. The bass performance you paid dearly for (it's expensive to get correct deep-bass reproduction) will work against you if the loudspeaker is put in a small room. For the same money, you could buy a superb minimonitor whose build cost was put into making the upper bass, midrange, and treble superlative. You win both ways with the minimonitor: your room won't be overloaded by bass, and the minimonitor will likely have much better soundstaging and tonal purity. There are other benefits: minimonitors, with their limited low-frequency extension, are less likely to annoy neighbors. You can thus listen to music louder without bothering anyone. Further, placement is much easier in small rooms.

Conversely, a minimonitor just won't fill a large room with sound. The sense of power, dynamic drive, deep-bass extension, and feeling of physical impact so satisfying in some music just doesn't happen with minimonitors. If you've got the room and the budget, a full-range, floorstanding loudspeaker is the best choice.

These are just a few examples of how the system you choose should be carefully tailored to your specific needs. We'll look at many more examples throughout this book.

Setting Your Budget

How much you should spend on a music-playback system depends on two things: your priorities in life and your financial means.

Let's take the priorities first. One person may consider a \$5000 stereo system an extravagance, yet not bat an eye at blowing \$15,000 on a two-week European vacation. Conversely, another person of similar means would find a \$15,000 vacation a waste of money when there is so much great hi-fi on the wish list. The first person probably considers music as merely a dispensable diversion while driving to work. To the second person, however, enjoying and appreciating music is a vital aspect of human existence. How much of your disposable income you should spend on a hi-fi system is a matter of how important music is to you, and only you can decide that. Owning a good hi-fi system will probably elevate music listening to a much higher priority in your life.

The second factor—your financial means—can often suggest a hi-fi budget. Surveys of readers of audiophile magazines suggest that about 10% to 15% of one's annual income is an average expenditure. In one recent survey, 25% of respondents owned systems costing less than \$10,000; 31% owned systems costing \$10,000 to \$35,000; 39% owned systems costing \$35,000 to \$75,000; and 5% of these magazine readers had more than \$75,000 invested in their audio systems. These are loose generalizations; high-end systems range in cost from less than \$1500 to more than \$500,000. And don't let these figures intimidate you: readers of audiophile magazines are dedicated enthusiasts who are not necessarily representative of all music lovers. You can put together a musically satisfying system for much less than the figures cited.

I strongly urge you, however, to establish a significant budget for your high-end system. The expenditure may seem high at the time, but you will be rewarded night after night and year after year with your favorite music wonderfully reproduced. A year or two from now, as you enjoy your system, the money spent will have been forgotten, but the pleasure will continue—a good music-playback system is a lasting and fulfilling investment. Moreover, if you buy a quality system now, you won't want to sell it or trade it in for something better later. It is sometimes false economy to "save" money on a less than adequate system. Do it right the first time.

There's another way of determining how much to spend for a hi-fi system: Find the level of quality you're happy with and let that set your budget. Visit your

dealer and have him play systems of various levels of quality. You may find yourself satisfied with a moderately priced system—or you may discover how good reproduced music can sound with the best equipment, and just *have* to have it.

When setting a budget for your music system, you must also consider whether the system will be used only for 2-channel music reproduction, or will be called on to reproduce movie soundtracks as part of a home-theater system. Maintaining high quality with home theater's seven speakers instead of two, and buying a multichannel surround-sound preamplifier rather than a 2-channel preamp, can significantly increase the amount of money you'll need to spend. There are, however, ways to add multichannel music reproduction and film-soundtrack playback without greatly escalating the system cost. We'll briefly touch on this subject in this chapter, and delve into it in more detail in [Chapter 11](#).

The Complete vs. the Incremental Purchase

After you've established a budget, you must decide which of the following three ways of buying a hi-fi is best for you:

- 1) Buy an entire system made up of the finest components.
- 2) Buy an entire system made up of components within a limited budget.
- 3) Buy just a few components now and add to the system as finances permit.

The first option doesn't require much thought—just a large bank account. The purchaser of this system needn't worry about budgets, upgrading, and adding components later. Other elements of buying a high-end system (I'll talk about these later) do apply to the cost-no-object audiophile: allocating the budget to specific components, dealing with the retailer, and home auditioning. But this sort of listener isn't under the same financial constraints that force the tough choices inherent in the other two options. Just choose a top-notch system and start enjoying it.

Most of us, however, don't enjoy such luxury, and must live within the set budgets of options 2 or 3. Option 2 is to spend the entire budget on a complete system now. In option 3, you'll spend your entire budget on just a few higher-quality components, and add the other pieces later as money becomes available. We'll call option 2 the *complete purchase*, and option 3 the *incremental purchase*. There are advantages and disadvantages to each approach.

Buying an entire system at once means that the overall budget must be spread among all the components that make up a hi-fi system. Consequently, you'll have less to spend on each component. You may not get the quality of components you'd hoped for, but the system will be complete and you can start enjoying it right away. This type of purchase is more suited to the music lover who doesn't want to think about equipment, but wants a system he can set up, forget, and use to enjoy music.

The listener more inclined to treat audio as a hobby, or who has her sights set on a more ambitious system, will buy just a few components now and add to the system as finances permit. This listener may spend the same amount of money as our first listener, but on just a pair of loudspeakers instead of on a whole system. She will use her old receiver or integrated amplifier, turntable, disc player, or portable music player until she can afford the electronics and sources she really

wants. Her system will be limited by the receiver's performance; she won't immediately get the benefit of her high-end loudspeakers. But when she *does* buy her dream electronics, she'll have a truly first-rate system. This approach requires more patience, commitment, and, in the long term, more money. But it is one way to end up with a superlative system.

The audiophile buying a system piece by piece can benefit from rapidly changing technology. She can start with components that don't change much over the years—power amplifiers, for example—and wait to buy those products likely to get better and cheaper over time: digital-to-analog converters (DACs), music servers, and so on.

Another benefit of adding components to your system one at a time is the ability to audition components in that system before buying. Rather than putting together a whole system in a store showroom based on your own auditioning or a salesperson's recommendations, the incremental buyer can carefully audition components and choose the best musical match for the rest of the system. This is a big advantage when assembling a hi-fi best suited to your musical tastes. This piece-by-piece approach is more for the hobbyist, and demands a deeper level of commitment to audio. It also requires a great deal more patience.

In short, if you want to get a good system, forget about the hardware, and just enjoy music, buy the entire system now. Given the same initial expenditure, you won't end up with the same-quality system as if you'd bought pieces slowly, but you'll be spending less money in the long run and can have the benefit of high-end music playback immediately. Moreover, you can go about your life without thinking about what piece of audio hardware to buy next. The complete purchase is recommended for the music lover who isn't—and doesn't want to be—an audio hobbyist. Many music lovers taking this approach, however, find themselves upgrading their systems piece by piece after discovering the rewards of owning a high-end audio system.

Conversely, the music lover on a budget who has set her sights on a more ambitious system, and who takes a more active role in audio equipment, will probably build a system gradually. This listener will more likely read product reviews in magazines, visit the dealer often, and take equipment home for evaluation. By doing so she'll become a better listener and a more critical audiophile, as well as develop a broader knowledge of audio equipment. Whichever approach you take, the information in the rest of this chapter—selecting a system suited to your listening, allocating the budget to specific components, and dealing with the retailer—applies equally well.

Value vs. Luxury Components

High-end audio components run the gamut from utilitarian-looking boxes to lavish, gold-plated, cost-no-object shrines. The packaging doesn't always reflect the quality of the electronics inside, but rather the manufacturer's product philosophy. Some companies try to offer the best sound for the least money by putting excellent electronics in inexpensive chassis. These are the so-called "value" products. Manufacturers of "luxury" components may put the same level of electronics in a lavish chassis with a 1"-thick front panel, lacquer-filled engraving, expensive metalwork, and custom machined input jacks or terminal posts.

A designer of *very* expensive electronics once told me that he could sell his products for *half* the price if he used a cheap chassis. He felt, however, that the level of design and execution in his electronics deserved no less than the ultimate in packaging.

Some buyers demand elegant appearance and a luxury feel; others merely want the best sound for the least money. To some music lovers, appearance and elegant packaging are secondary to sound quality; they don't care what it looks like so long as it sounds good. Conversely, some audiophiles are willing to pay for gorgeous cosmetics, battleship build quality, and all the trimmings that make some products exude elegance and luxury. There's an undeniable pride of ownership that accompanies the finest-built audio components.

When choosing high-end components, match your needs with the manufacturer's product philosophy. That way, you won't waste money on thick faceplates you don't care about, or spend money on a product that doesn't do justice to your home decor.

Another type of manufacturer puts mediocre or even poor electronics in a fancy, eye-catching package. Their market is the less musically sophisticated buyer who chooses on the basis of appearance and status (or the company's past reputation) rather than sound quality. These so-called "boutique" brands are not high-end, however, and should be avoided.

Allocating Your Budget to Specific Components

There are no set rules for how much of your total budget you should spend on each component in your system. Allocating your budget between components depends greatly on which components you choose, and your overall audio philosophy. Mass-market mid-fi magazines have been telling their readers for years to spend most of a hi-fi budget on the loudspeakers because they ultimately produce the sound. This thinking also suggests that all amplifiers and digital sources sound alike; why waste money on expensive amplifiers and DACs?

The high-end listener makes different assumptions about music reproduction. A fundamental tenet of high-end audio holds that if the signal isn't good at the beginning of the reproduction chain, nothing downstream can ever improve it. In fact, the signal will only be degraded by any product it flows through. High-end audio equipment simply minimizes that degradation. If your DAC or phono cartridge is bright, hard, and unmusical, the final sound will be bright, hard, and unmusical. Similarly, the total system's performance is limited by the resolution of the worst component in the signal path. You may have superb loudspeakers and an excellent turntable and cartridge, but they'll be wasted with a low-resolution preamp in the signal chain.

Quality matching between components is essential to getting the most sound for your budget. High-quality loudspeakers at the end of a chain containing a bad-sounding component can even make the system sound worse than lower-quality loudspeakers: The high-resolution loudspeakers reveal all the imperfections of the electronics upstream of them. This situation has been likened to having a large picture window in your home. If the view is of the Northern California coastline, you want that window to be as clean and transparent as possible. But if the window overlooks a garbage dump, you'd prefer that it somewhat obscure the view.

In the course of my reviewing career, I've listened to \$400 loudspeakers driven by \$30,000 worth of electronics, and \$158,000 loudspeakers driven by budget integrated amplifiers. I can state categorically that the electronics and source components are every bit as important as the loudspeakers. Although the loudspeakers significantly influence the overall sound, high-quality source components (turntable and DAC), good electronics (preamplifier and power amplifier), and excellent cables are essential to realizing a musical high-end system.

Since the first four editions of this book, however, I've had experiences with very good-sounding and affordable integrated amplifiers driving moderately expensive loudspeakers, and the results have been excellent. As explained in [Chapters 5 and 6](#), loudspeaker sensitivity (how loudly the speaker will play for a given amount of amplifier power) greatly affects how powerful an amplifier you need to achieve a satisfying volume. High-sensitivity speakers need very little amplifier power. And because amplifier power is costly, it follows that a system with high-sensitivity speakers can be driven by lower-cost amplification—*provided that the amplification is of high quality*. Fortunately, manufacturers have responded to the need for relatively inexpensive high-performance amplification by designing integrated amplifiers with outstanding sound quality, but with lower output powers. By putting the preamplifier and power amplifier in one housing and cutting back on power output, manufacturers can put their high-end circuits in lower-priced products. The trick is to find those bargain integrated amplifiers that deliver truly high-end sound, and mate them with loudspeakers that not only have the appropriate sensitivity, but are also a good musical match. This approach will get you the best sound for the least money. If, however, cost is secondary to sound quality—that is, you're willing to spend more for an improvement in sound—then buy the best separate preamplifier and power amplifier you can find.

For the following exercise, I assembled an imaginary system of the components I'd choose if my audio budget totaled \$11,000. This hypothetical system follows a traditional audiophile approach. Here are the costs per item:

Preamplifier	\$2000
Power amplifier	\$2000
Digital source	\$1300
Loudspeakers	\$4000
Interconnects and cables	\$700
Power conditioner and AC cords	\$1000
Total	\$11,000

As you can see, loudspeakers consumed about 36% of the budget, the digital source took up another 12%, and the preamp and power amplifier each received 18%. About 9% was allocated to an AC power conditioner and AC cords, and the remaining 6% was spent on interconnects and cables. You may be surprised to see nearly 10% of the system budget devoted to an AC power conditioner and

specialty AC cords, but AC conditioning is essential to good sound and offers a high sonic return on investment. These numbers and percentages aren't cast in stone, but they're a good starting point in allocating your budget. If you wanted to include a turntable, tonearm, and cartridge, the budget for the other components would have to be reduced.

Another approach with this budget would be to buy an inexpensive DAC (\$500) and put the \$800 saved into better loudspeakers or electronics. Then, as finances permitted, you could upgrade the DAC, a component category that gets better and less expensive over time (see [Chapter 7](#)). You'll have music in the meantime, and end up with a better-sounding system in the long run. You could also choose a preamplifier with an integral DAC, a feature included in most entry-level and mid-priced models. This approach, however, precludes upgrading just the DAC at a later date.

The 36% figure for loudspeakers is very flexible. Keep in mind that there's an important threshold in loudspeaker performance at about \$2500: Loudspeakers costing a little more than \$2500 are often disproportionately better than those costing a little less than \$2500, so you may want to adjust your budget allocations to cross this threshold. As described in [Chapter 6](#), many moderately priced loudspeakers outperform much more expensive models. Use [Chapter 6's](#) guidelines on choosing loudspeakers to get the most performance for your loudspeaker dollar.

Following the earlier discussion of matching a superb but low-powered integrated amplifier with high-sensitivity speakers, here's another example of how I might allocate an \$11,000 budget:

Integrated amplifier with DAC	\$4000
Loudspeakers	\$5500
Interconnects and cables	\$500
Power conditioner and AC cords	\$1000
Total	\$11,000

Again, the key to putting so much of the budget into loudspeakers is extremely careful matching of the amplifier's power output to the loudspeaker's sensitivity (and impedance curve, explained in [Chapter 6](#)), along with finding those few integrated amplifiers that deliver the musicality of expensive separates, but simply have lower output powers. Here's an extreme case: I lived with a system for about a month (during a product review) that included \$11,000

loudspeakers driven by a \$1500 integrated amplifier, and the result was musical magic. It takes a lot of searching to find these synergistic combinations—or a great dealer who has discovered these ideal matches for you. I must stress that this approach works only with certain components, and is useful for getting the best sound for the least money. It is not the ideal strategy when the best possible sound is your goal.

Here's another sample budget, this one based on a maximum expenditure of \$2000:

Integrated amplifier with DAC	\$1000
Loudspeakers	\$750
Interconnects and cables	\$250
Total	\$2000

Again, I selected components that experience suggested would be a good match, and tallied the percentages after choosing the components. Interestingly, the breakdown was similar to that in the first example: 37% on loudspeakers and about 10% on interconnects and cables.

I've heard systems at this price level that are absolutely stunning musically. When carefully chosen and set up, a \$2000 high-end system can achieve the essence of what high-quality music reproduction is all about—communicating the musical message. I've even heard a system with a total list price of only \$850 that was musical and enjoyable. The point isn't how much you spend on a hi-fi, but how carefully you can choose components to make a satisfying system within your budget.

These guidelines are for a dedicated 2-channel music system, not a multichannel music and home-theater system. The different requirements of home theater suggest a somewhat different budget allocation. If most of your listening time with a multichannel system is devoted to music, the budget guidelines are very similar to those described earlier, but with a shift of money from electronics to loudspeakers. That is, you should spend a bit more of your total budget on speakers. That's because home theater and multichannel music require five or seven loudspeakers (plus a subwoofer), and spreading out a fixed amount of money over five or seven speakers rather than two inevitably results in compromised performance. If music is your primary consideration, put most of your speaker budget into the left and right speakers, with less devoted to the center, surround speakers, and subwoofer. Those readers who spend more time watching movies should spend even less on electronics and source and more on

speakers. That's because the center-channel loudspeaker plays an important role in the home-theater experience (it reproduces nearly all the dialogue), and choosing a quality model is of paramount importance.

In my experience, the sonic differences among amplifiers, preamplifiers, and source components are much less pronounced when we watch movies than when we listen to music. So much of our attention is taken away from the aural experience by the overwhelming visual sensory input. I've used audio/video receivers in my system (during product reviews) and found most of them lacking when reproducing music, but more than adequate for film-soundtrack playback.

Whether choosing a dedicated 2-channel music system or a combined music and home-theater package, you should save some of your budget for an AC power conditioner and accessories. I advise against buying a power conditioner and accessories when you buy the system unless you have confidence in the conditioner from reading credible reviews. A good approach is to take the audio system home, get it set up and optimized, *then* add a power conditioner and start experimenting with accessories. Here's why: AC conditioners don't always make an improvement. In fact, they can degrade the sound. There are many variables with AC power conditioners, including the quality of AC from your wall, the method of AC conditioning, and the number and nature of the components plugged into the conditioner. It is therefore best to try the conditioner at home before buying.

There's another good reason for adding an AC line conditioner later: By getting to know how your system sounds *without* an AC conditioner, you'll be better able to judge if the conditioner is an improvement. Remember that a change in sound isn't always for the better. The same logic holds true for accessories such as cones, feet, and tube dampers: You'll be in a much better position to judge their effectiveness—or lack of it—by knowing your system intimately *before* installing accessories. Set aside some of your budget—perhaps a few hundred dollars—for accessories. If they don't make a difference, you've now got a few hundred dollars to spend on records or CDs. (A full survey of accessories is included in [Chapter 15](#).)

Upgrading a Single Component

Many audiophiles gradually improve their systems by replacing one component at a time. The trick to getting the most improvement for the money is to replace the least-good component in your system. A poor-sounding preamp won't let you hear how good your music server is, for example. Conversely, a very clean and transparent preamp used with a grainy and hard digital source will let you hear only how grainy and hard the digital source is. The system should be of similar quality throughout. If there's a quality mismatch, however, it should be in favor of high-quality source components.

Determining which component to upgrade can be difficult. This is where a good high-end audio retailer's advice is invaluable—he can often pinpoint which component you should consider upgrading first. Another way is to borrow components from a friend and see how they sound in your system. Listen for which component makes the biggest improvement in the sound. Finally, you can get an idea of the relative quality of your components by carefully reading the high-end audio magazines, particularly when they recommend specific components.

In [Chapter 1](#), I likened listening to music through a playback system to looking at the Grand Canyon through a series of panes of glass. Each pane distorts the image in a different way. The fewer and more transparent the panes are, the clearer the view, and the closer the connection to the direct experience.

Think of each component of a high-end audio system as one of those panes of glass. Some of the panes are relatively clear, while others tend to have an ugly coating that distorts the image. The pane closest to you is the loudspeaker; the next closest pane is the power amplifier; next comes the preamp; and the last pane is the signal source (disc player, music server, turntable). Your view of the music—the system's overall transparency—is the sum of the panes. You may have a few very transparent panes, but the view is still clouded by the dirtiest, most colored panes. This idea is shown graphically on [page 3](#).

The key to upgrading a hi-fi system is to get rid of those panes—those components that most degrade the music performance—and replace them with clearer, cleaner ones. This technique gives you the biggest improvement in sound quality for the money spent.

Conversely, putting closest to you a very transparent pane—the loudspeaker—only reveals in greater detail what's wrong with the power amplifier,

preamplifier, and source components. A high-resolution loudspeaker at the end of a mediocre electronics chain can actually sound worse than the same system with a lower-quality loudspeaker.

Following this logic, we can see that a hi-fi system can never be any better than its source components. If the first pane of glass—the source component—is ugly, colored, and distorts the image, the result will be an ugly, colored, and distorted view. Upgrading a hi-fi system is a fine art that requires that you not only target the weakest link, but also that you select each new component very carefully; your next upgrade will unmask every other component's intrinsic sonic personality. It's also essential that the upgraded component is a demonstrable step forward in transparency and resolution, not merely one that offers a different set of virtues and flaws that tilts your system toward a sound that is different but no better. The latter is the path to years of "upgrades" that result in no real increase in musical realism.

As you upgrade your system, you can start to see that other panes you thought were transparent actually have some flaws you couldn't detect before. The next upgrade step is to identify and replace what is *now* the weakest component in the system. This can easily become an ongoing process.

Unfortunately, as the level of quality of your playback system rises, your standard of what constitutes good performance rises with it. You may become ever more critical, upgrading component after component in the search for musical satisfaction. This pursuit can become an addiction and ultimately *diminish* your ability to enjoy music. The next chapter includes an editorial I wrote for *Stereophile* magazine examining this subject.

How to Read Magazine Reviews

I've deliberately avoided recommending specific products in this book. By the time you would have read the recommendations, the products would likely have been updated or discontinued. The best source for advice on components is product reviews in high-end audio magazines. Because most of these magazines are published monthly, they can stay on top of new products and offer up-to-date buying advice. Many magazine reviewers are highly skilled listeners, technically competent, and uncompromising in their willingness to report truthfully about audio products. There's a saying in journalism that epitomizes the ethic of many high-end product reviewers: "without fear or favor." The magazine should neither fear the manufacturer when publishing a negative review, nor expect favor for publishing a positive one. Instead, the competent review provides unbiased and educated opinion about the sound, build quality, and value of individual products. Because high-end reviewers hear lots of products under good conditions, they are in an ideal position to assess the relative merits and drawbacks and report their informed opinions. The best reviewers have a combination of good ears, honesty, and technical competence.

You'll notice a big difference between reviews in high-end magazines and reviews in the so-called "mainstream slicks." The mass-market, mainstream magazines are advertiser-driven; their constituents are their advertisers. Conversely, high-end magazines are usually reader-driven; the magazines' goal is to serve their readers, not their advertisers. Consequently, high-end magazines often publish negative reviews, while mass-market magazines generally do not. Moreover, high-end magazines are much more discriminating about what makes a good component that can be recommended to readers. Mass-market magazines cater to the average person in the street, who, they believe (mistakenly, in my view), doesn't care about aspects of music reproduction important to the audiophile. The high-end product review is thus not only more honest, but much more discriminating in determining what is a worthy product. If you're reading a hi-fi magazine that never criticizes products, beware. Not all audio components are worth buying; therefore not all magazine reviews should conclude with a recommendation. Also beware of so-called "consumer" publications that regard an audio system as an appliance, not as a vehicle for communicating an art form. In their view, the improvements in performance gained by spending more are not worth the money, so they recommend the least expensive products with the most

features. They take a similar approach to reviewing cars: why spend lots of money on a depreciating asset when all cars perform an identical function—moving you from point A to point B? It is simply not within their ethos to value spirited acceleration, tight handling, a luxurious interior, or the overall feel of the driving experience.

As someone who makes his living reviewing high-end audio products, I'll let you in on a few secrets that can help you use magazine reviews to your advantage. First, associate the review you're reading with the reviewer. Before reading the review, look at the byline and keep in mind who's writing the review. This way, you'll quickly learn different reviewers' tastes in music and equipment. Seek the guidance of reviewers with musical sensibilities similar to your own.

Second, don't assign equal weight to all reviews or reviewers. Audio reviewing is like any other field of expertise—there are many different levels of competence. Some reviewers have practiced their craft for decades, while others are newcomers who lack the seasoned veteran's commitment to the profession. Consider the reviewer's reputation, experience, and track record when giving a review credence. Also consider the reviewer's standards for what makes a good product. Just as there's a big difference between a movie recommendation given by Joel Siegel (*Good Morning America*) and Andrew Sarris (*The Village Voice*), so too are there differences in product recommendations from high-end reviewers. The rise of the Internet has allowed anyone with a computer to publish audio-equipment reviews. A few of the online magazines employ skilled listeners and adhere to high journalistic standards; others are merely vehicles for audiophiles seeking free use of review equipment, self-aggrandizement, or both.

Finally, listen to the products yourself. If two components are compared in a review, compare those products to hear if your perceptions match the reviewer's. Even if you aren't in the market for the product under review, listening to the components will sharpen your skills and put the reviewer's value judgments in perspective.

A common mistake among audiophiles looking for guidance is to select a component on the basis of a rave review without fully auditioning the product for themselves. The review should be a starting point, not some final judgment of component quality. It's much easier to buy a product because a particular reviewer liked it than it is to research the product's merits and shortcomings for yourself. Buying products solely on the basis of a review is fraught with danger. Never forget that a review is nothing more than one person's opinion, however informed and educated that opinion might be. Moreover, if the reviewer's tastes differ from yours, you may end up with a component you don't like. We all have

different priorities in judging reproduced sound quality; what the reviewer values most—soundstaging, for example—may be lower in your sonic hierarchy. Your priorities are the most important consideration when choosing music-playback components. Trust your own ears.

Sharp-eyed readers can find out which products a reviewer thinks are really special. Most reviewers list the reference system with which they audition new products. When a particular product shows up in the reference system in review after review, you can be sure that the writer thinks that product is especially good.

Some magazine readers view all reviews with suspicion because the reviewer enjoyed the free loan of the product. That is, they believe there's a quid pro quo that guarantees the manufacturer a positive review because the reviewer got to use the product for several months. This belief is patently false. Reviewers have access to virtually any products they want, and have no allegiance to the manufacturer who lent it to them. This may be hard to believe from a reader's standpoint, but after receiving, unpacking, setting up, and learning about hundreds of products, getting new equipment often becomes a chore rather than cause for excitement.

I'd also like to address the issue of long-term equipment loans. Most reviewers have reference systems composed of products they don't own. The manufacturers or distributors of those products choose to allow the reviewer to keep the product for many months or even a year, giving the product exposure in issue after issue. Reviewers need reference-quality equipment with which to judge the product under evaluation—reference-quality equipment they can't afford to buy. The long-term loan is mutually beneficial: the reviewer has access to the best tools available; the manufacturer enjoys the prestige—and, presumably, increased sales—conferred by the reviewer's repeated mentions; and the reader gets the most accurate descriptions of the products reviewed within the context of the reference system. Moreover, because so many manufacturers see the value in long-term loans, reviewers have access to virtually any products they want. This removes the putative pressure to give positive reviews only to those manufacturers who agree to long-term loans. If a reviewer keeps a product in his system, it's because that product sounds good.

It would be a violation of journalistic ethics, however, if a deal were struck between manufacturer and reviewer that allowed the reviewer to keep the product as his own at the conclusion of the review. In my 25 years' experience as a reviewer and magazine editor, I have never been approached by a manufacturer in an untoward way, have never been aware of such behavior (at least on the

magazines I've worked for), and would not tolerate any such ethical lapses by a writer on our staff.

A final word about product reviews: If a component you've already purchased gets a negative review, don't let that end your enjoyment of it. You enjoyed music through that product before the review appeared; why should another person's opinion diminish your own musical satisfaction? Positive or negative, magazine reviews carry far too much weight in readers' minds and in the marketplace. (Remember, you heard this from a professional reviewer.)

In sum, reviews can be very useful, provided you:

- Get to know individual reviewers' sonic and musical priorities. Find a reviewer on your sonic and musical wavelength, then trust his or her opinions.
- Compare your impressions of products to the reviewer's impressions. This will not only give you a feel for the reviewer's tastes and skill, but the exercise will make you a better listener.
- Don't assign the same weight to all reviews. Consider the reviewer's reputation and experience in the field. How many similar products has the reviewer auditioned? If a reviewer has heard virtually every serious phono cartridge, for example, his or her opinion will be worth more than that of the reviewer who has heard only a few models.
- Don't buy—or summarily reject—products solely on the basis of a review. Use product recommendations as a starting point for your own auditioning. Listen to the product yourself to decide if that product is for you. Let *your* ears decide.

System Matching

It is a truism of high-end audio that an inexpensive system can often outperform a more costly and ambitious rig. I've heard modest systems costing, say, \$1500 that are more musically involving than \$50,000 behemoths. Why?

Part of the answer is that some well-designed budget components sound better than ill-conceived or poorly executed esoteric products. But the most important factor in a playback system's musicality is *system matching*. System matching is the art of putting together components that complement each other sonically so that the overall result is a musicality beyond what each of the components could achieve if combined with less compatible products. The concept of synergy—that the whole is greater than the sum of the parts—is very important in creating the best-sounding system for the least money.

System matching is the last step in choosing an audio system. You should have first defined the system in terms of your individual needs, set your budget, and established a relationship with a local specialty audio retailer (if there's a dealer in your immediate area). After you've narrowed down your choices, which products you select will greatly depend on system matching.

Knowing which components work best with which other components is best learned by listening to a wide range of equipment. Many of you don't have the time—or access to many diverse components—to find out for yourselves what equipment works best with what other equipment. Consequently, you must rely on experts for general guidance, and on your own ears for choosing specific equipment combinations.

The two best sources for this information are magazine reviews and your local dealer. Your dealer will have the greatest knowledge about products he carries, and can make system-matching recommendations based on his experience in assembling systems for his customers. Your dealer will likely have auditioned the products he sells in a variety of configurations; you can benefit from his experience by following his system-matching recommendations. With the unfortunate decline in the number of audio dealers, a few catalog-based retailers have emerged that offer a wide selection of products along with expert advice.

The other source of system-matching tips is magazine reviews. Product reviews published in reputable magazines will often name the associated equipment used in evaluating the product under review. The reviewer will sometimes describe his or her experiences with other equipment not directly part

of the review. For example, a loudspeaker review may include a report on how the loudspeaker sounded when driven by three or four different power amplifiers. The sonic characteristics of each combination will be described, giving the reader an insight into which amplifier was the best match for that loudspeaker. More important, however, the sonic descriptions and judgments expressed can suggest the *type* of amplifier best suited to that loudspeaker. By *type* I mean both technical performance (tubed vs. transistor, power output, output impedance, etc.) and general sonic characteristics (hard treble, forward presentation, well-controlled bass, etc.).

Let's say the reviewer drove the loudspeakers with four amplifiers: a low-powered but sweet-sounding integrated amplifier, a high-output-impedance tubed design, a medium- to high-powered inexpensive solid-state unit, and a massive solid-state amplifier that requires two people to lug it into the listening room. The reviewer reports that the integrated amplifier just didn't have enough power to produce sufficient volume, and that the sound lacked dynamics. The high-output-impedance tubed amplifier was mushy in the bass and had a reduced sense of pace and rhythm. The inexpensive solid-state amplifier had terrific bass control, but its forward presentation and grainy treble made it less than ideal with the loudspeaker under review. (All of these terms—*forward*, *grainy*, *pace*, etc.—are defined in the next chapter.) Finally, the reviewer concludes that the solid-state behemoth is the only amplifier suitable for this particular loudspeaker.

This doesn't mean that the most expensive amplifier will work best with *every* loudspeaker. I can think of loudspeakers for which this scenario would be completely different. Another loudspeaker would sound just fine with the integrated amplifier, suggesting that using the huge solid-state unit would be overkill. However, if the loudspeaker was a little *tizzy* in the treble and lean in the midbass, the tubed amplifier would tend to ameliorate these tendencies. Finally, if the loudspeaker wasn't that sensitive to treble grain, but needed to be driven by an amplifier with control and authority in the bass, the inexpensive solid-state unit would be a good choice—and the most cost-effective.

These reports of system matching can extend beyond the specific products reported on in the review. A fairly good idea of which type of sonic and technical performance is best suited to a particular product can be gained from a careful reading of product reviews. For example, you may conclude that a particular loudspeaker needs to be driven by a large, high-current amplifier. This knowledge can then point you in the right direction for equipment to audition yourself: you can rule out low-powered designs.

By reading magazine reviews, following your dealer's advice, and listening to combinations of products for yourself, you can assemble a well-matched system that squeezes the highest musical performance from your hi-fi budget.

Dos and Don'ts of Selecting Components

Some audiophiles are tempted to buy certain products for the wrong reasons. For example, many high-end products are marketed on the basis of some technical aspect of their design. A power amplifier may, for example, be touted as having “over 200,000 microfarads (μF) of filter capacitance,” “32 high-current output devices,” and a “discrete JFET input stage.” While these may be laudable attributes, they don't guarantee that the amplifier will produce good sound. Don't be swayed by technical claims—listen to the product for yourself.

Just as you shouldn't make a purchasing decision based on specifications, neither should you base your decision solely on brand name. Many high-end manufacturers with solid reputations sometimes produce mediocre-sounding products. A high-end marque doesn't necessarily mean high-end sound. Again, let your ears be your guide. I'm often pleasantly surprised to find moderately priced products that sound as good—or very nearly as good—as products costing two or three times as much.

When choosing components, however, you should consider the company's longevity, reputation for build quality, customer service record, and product reliability. High-end manufacturers run the gamut from one-man garage operations to companies with hundreds of employees and advanced design and manufacturing facilities. The garage operation may produce good-sounding products, but may not be in business next year. This will not only make it hard to get service, but also greatly lowers the product's resale value.

High-end manufacturers also have very different policies regarding service. Some repair their products grudgingly, and/or charge high fees for fixing products out of warranty. Others bend over backward to keep their valued customers happy. In fact, some high-end audio companies go to extraordinary lengths to please their customers. One amplifier manufacturer who received an out-of-warranty product for repair not only fixed the amplifier free of charge, but replaced the customer's scratched faceplate at no cost! It pays in the long run to do business with manufacturers who have reputations for good customer service.

Another factor to consider before laying down your hard-earned cash is how long the product has been on the market. Without warning, manufacturers often discontinue products and replace them with new ones, or update a product to “Mark II” status. When this happens, the value of the older product drops immediately. If you know an excellent product is about to be discontinued, you

can often buy the floor sample at a discount. This is a good way of saving money, provided the discount is significant. You end up with a lower price, plus all the service and support inherent in buying from an authorized and reputable dealer rather than a private party.

The best source of advance information on new products and what's about to be discontinued are reports in audio magazines and their websites from the many consumer and trade shows held around the world.

Your Relationship with the Retailer

If, in the past, you've bought audio equipment only from mass-market retailers, you should expect to have an entirely different sort of relationship with a high-end dealer. The good specialty audio retailer doesn't just "move" boxes of electronics; he provides you with the satisfaction of great-sounding music in your home. More than just an equipment dealer, he's usually a dedicated audio and music enthusiast himself—he *knows* his products, and is often the best person to advise you on selecting equipment and system setup. The great dealers search the entire audio landscape for the very best brands and are extremely selective in choosing which lines to carry.

Consider the very different relationships between seller and buyer in the following scenarios. In the first, a used-car dealer in downtown Los Angeles is trying to sell a car to someone from out of town. The seller has only one shot at the buyer, and he intends to make the most of it. He doesn't care about return business, the customer's long-term satisfaction with the purchase, or what the customer will tell his friends about the dealer. It will be an adversarial relationship from start to finish.

Then consider a new-car dealer in Great Falls, Montana, selling a car to another Great Falls resident. For this dealer, return business is vital to his survival. So is customer satisfaction, quality service, providing expert advice on models and options, finding exactly the right car for the particular buyer, and giving the customer an occasional ride to work when he drops off his car for service. He knows his customers by name, and has developed mutually beneficial, long-term relationships with them.

If buying a mass-market hi-fi system is like negotiating with a used-car dealer in downtown L.A., selecting a high-end music system should be made within a relationship similar to the one enjoyed between the Montana dealer and his customers.

Take the time to establish a relationship with your local dealer. Make friends with him—it'll pay off in the long run. Get to know a particular salesperson and, if possible, the store's owner. Tell them your musical tastes, needs, lifestyle, and budget—then let them offer equipment suggestions. They know their products best, and can offer specific component recommendations. The good stores will regard you as a valued, long-term customer, not someone with whom they have

one shot at making a sale. Don't shop just for equipment—shop for the retailer with the greatest honesty and competence.

Keep in mind, however, that dealers will naturally favor the brands they carry. Be suspicious of dealers who badmouth competing brands that have earned good reputations in the high-end audio press. The best starting point in assembling your system is a healthy mix of your dealer's recommendations and unbiased, competent magazine reviews.

The high-end retailing business is very different from the mass-market merchandising of the low-quality "home entertainment" products sold in appliance emporiums. The specialty retailer's annual turnover is vastly lower than that of the mid-fi store down the street. Consequently, the specialty retailer's profit margin must be larger for him to stay in business. Don't expect him to offer huge discounts and price cuts on equipment, to take a sale from a mid-fi store. Because the high-end dealer offers so much more than just pushing a box over the counter, his prices just can't be competitive. Instead, you should be prepared to pay full list price—or very close to it.

Here's why. After paying his employees, rent, lights, heat, insurance, advertising, and a host of other expenses, the specialty audio retailer can expect to put in his pocket about five cents out of every dollar spent in his store. Now, if he discounts his price by even as little as 5%, he is essentially working for free. If the dealer offers a discount or marks down demonstration or discontinued units, you should take advantage of these opportunities. But don't expect the dealer to discount; he deserves the full margin provided by the product's suggested retail price.

In return for paying full price, however, you should receive a level of service and professionalism second to none. Expect the best from your dealer. Spend as much time as you feel is necessary auditioning components in the showroom before you buy. Listen to components at home in your own system before making a purchasing decision. Ask the retailer to set up your system for you. Exploit the dealer's wide knowledge of which components are best for the money. Use his knowledge of system matching to get the best sound possible on a given budget. And if one of your components needs repair, don't be afraid to ask for a loaner until yours is fixed. The dealer should bend over backward to accommodate your needs.

If you give the high-end dealer your loyalty, you can expect this red-carpet treatment. This relationship can be undermined, however, if, to save a few dollars, you buy from a competitor or by mail a product that your dealer also sells. If the

product purchased elsewhere doesn't sound good in your system, don't expect your local dealer to help you out. Further, don't abuse the home-audition privilege. Take home only those products you're seriously considering buying. If the dealer let everyone take equipment home for an audition, he'd have nothing in the store to demonstrate. The home audition should be used to confirm that you've selected the right component through store auditioning, magazine reviews, and the dealer's recommendations. The higher price charged by the dealer may seem hard to justify at first, but in the long run you'll benefit from his expertise and commitment to you as a customer.

If you don't live close to any high-end dealers, there are several very good mailorder companies that offer excellent audio advice over the phone. They provide as much service as possible by phone, including money-back guarantees, product exchanges, and component-matching suggestions. You can't audition components in a store, but you can often listen to them in your system and get a refund if the product isn't what you'd hoped it would be.

In short, if you treat your dealer right, you can expect his full expertise and commitment to getting you the best sound possible. There's absolutely no substitute for a skilled dealer's services and commitment to your satisfaction.

Used Equipment

A used audio component often sells for half its original list price, making used gear a tempting alternative. The lower prices on used high-end components provide an opportunity to get a high-quality system for the same budget as a less ambitious new system. Moreover, buying used gear lets you audition many components at length. If you find a product you like, keep it. If you don't like the sound of your used purchase, you can often sell it for no less than what you paid.

There are two ways of buying used equipment: from a dealer, or from a private party, usually over an Internet audio-trading site. A retailer may charge a little more for used products, but often offers a short warranty (60 or 90 days), and sometimes exchange privileges. Buying used gear from a reputable retailer is a lot less risky than dealing with a private party (unless you're buying from a friend).

The audiophile inclined to buy used equipment can often get great deals. Some audiophiles simply *must* own the latest and greatest product, no matter the cost. They'll buy a state-of-the-art component one year, only to sell it the next to acquire the current top of the line. These audiophiles generally take good care of the equipment and sell it at bargain-basement prices. If you can find such a person, have him put your name at the top of his calling list when he's ready to sell. You can end up with a superb system for a fraction of its original selling price.

A few pitfalls await the buyer of used equipment, though the disadvantages listed below apply mostly to buying from a private party rather than from a reputable dealer. First, there are no assurances that the component is working properly; the product could have a defect not apparent from a cursory examination. Second, the used product could be so outdated that its performance falls far short of new gear selling for the same price—or less—than the used component. This is especially true of digital products—DACs, disc players, and music servers. Third, a used product carries no warranty; you'll have to pay for any repair work. Finally, you must ask why the person is selling the used equipment. All too often, a music lover doesn't do his homework and buys a product that doesn't satisfy musically or work well with the rest of his system. If you see lots of people selling the same product, beware—it's a sign that the product has a fundamental musical flaw. Finally, buying used equipment from a private party eliminates everything that makes your local specialty retailer such

an asset when selecting equipment: You don't get the dealer's expert opinions, home audition, trade or upgrade policy, dealer setup, warranty, dealer service, loaner units, or any of the other benefits you get from buying new at a dealer.

Approach used components with caution; they can be a windfall—or a nightmare.

Product Upgrades

Many manufacturers improve their products and offer customers the option of upgrading their components to current performance. This is the “Mark II” designation on some components. The dealer can usually handle sending your component back to the factory. Some manufacturers prefer to deal with the customer directly, saving the dealer markup and keeping the upgrade price lower.

I have ambivalent feelings about manufacturer upgrades. Some upgrade programs provide lasting value to a company’s customers. Enjoying the benefits of a company’s advances without selling a component at a loss and buying a new one can be a wonderful bonus. Conversely, some manufacturers think of an upgrade program as a profit center, charging large amounts for even minor improvements. Consider a company’s track record when evaluating potential future upgrades. For products that are upgradeable via software updates, ask how much these revisions will cost. Some manufacturers charge nothing, or a nominal fee, for updated software.

Component Selection Summary

When choosing a high-end system or component, follow these ten guidelines:

- 1) Establish your budget. Buy a component or system you'll be happy with in the long run, not one that will "do" for now. Do it right the first time.
- 2) Be an informed consumer—learn all you can about high-end audio. Study magazine reviews, visit your local specialty retailer, and read the rest of this book. Do your homework.
- 3) Develop a relationship with your dealer. He can be the best source of information in choosing components and assembling a system.
- 4) Find components that work synergistically. Again, your dealer knows his products and can offer suggestions.
- 5) Select products based on their musical qualities—not technical performance, favorable reviews, specifications, price, or brand name.
- 6) Choose carefully; many lower-priced components can outperform higher-priced ones. Take your time and maintain high standards—there are some great bargains out there.
- 7) Buy products from companies with good reputations for value, customer service, and reliability. Also, match the company's product philosophy (e.g., cost-no-object vs. best value for the money) to your needs.
- 8) When possible, listen to prospective components in your system at home before buying.
- 9) Follow the setup guidelines in [Chapters 14 and 15](#) to get the most from your system. Enlist the aid of your dealer in system setup.
- 10) When upgrading a single component, be sure that the new component is truly more transparent and resolving than the component it replaces, and not just a lateral step that substitutes one set of sonic trade-offs for a different set of trade-offs.

11) Add accessories *after* your system is set up.

If you read the rest of this book, subscribe to one or more reputable high-end magazines, and follow these guidelines, you'll be well on your way to making the best purchasing decisions—and having high-quality music reproduction in your home.

One last piece of advice: After you get your system set up, forget about the hardware. It's time to start enjoying music.

3 Becoming a Better Listener

Critical listening—the practice of evaluating the quality of audio equipment by careful analytical listening—is very different from listening for pleasure. The goal isn't to enjoy the musical experience, but to determine if a system or component sounds good or bad, and which *specific characteristics* of the sound make it good or bad. You want to critically examine what you're hearing so that you can form judgments about the reproduced sound. You can then use this information to evaluate and choose components, and to fine-tune a system for greater musical enjoyment.

Evaluating audio equipment by ear is essential—today's technical measurements simply aren't advanced enough to characterize the musical performance of audio products. The human hearing mechanism is vastly more sensitive and complex than the most sophisticated test equipment now available. Though technical performance is a valid consideration when choosing equipment, the final arbiter of good sound should always be the ear.

Moreover, the musical significance of sonic differences between components can only be judged subjectively. This is best expressed by Michael Polanyi in his book *Personal Knowledge*:

Whenever we find connoisseurship operating within science or technology we may assume it persists only because it has not been possible to replace it by a measurable grading.

Because connoisseurship is used in the evaluation of audio products, the process is more an aesthetic endeavor than a purely technical one. Good technical performance can contribute to high-quality musical performance, but it doesn't tell you what you really want to know: how well the product communicates the musical message. To find that out, you must listen. I have auditioned hundreds of audio products for review and measured their technical performance in a test laboratory. My experience overwhelmingly indicates that much more about the

quality of an audio component can be learned in the listening room than in the test lab.

Many newcomers to high-performance music reproduction—and even a small fringe group of experienced audiophiles—question the need for listening to evaluate products. They believe that measurements can tell them everything they need to know about a product’s performance. And since these measurements are purely “objective,” why interject human subjectivity through critical listening?

The answer is that the common measurements in use today were created decades ago as design tools, not as descriptors of sound quality. The test data generated by a typical mix of audio measurements were never meant to be a representation of musical reality, only a rough guide when designing. For example, an amplifier circuit that had 1% harmonic distortion was probably better than one with 10% harmonic distortion. It doesn’t follow that a harmonic-distortion specification in any way describes the sound of that amplifier.

A second problem is that audio test-bench measurements attempt to quantify a variety of two-dimensional phenomena: how much distortion the product introduces at a given signal level, its frequency response, noise level, and other factors. But music listening is a *three-dimensional* experience that is much more complex than any set of numbers can hope to quantify. How can you reduce to a series of mathematical symbols the ability of one power amplifier, and not another, to make the hair on your arms stand up? Or the feeling that a vocalist is singing directly *to you*?

I once read that a motorcycle manufacturer attempted to quantify the “feel” of prototype motorcycles by attaching hundreds of sensors at many positions around the motorcycle and putting the motorcycle on a dynamometer in a wind tunnel. The large amount of data collected was then processed to guide the engineers in designing improvements in the motorcycle’s handling and overall riding experience. After investing a significant amount of money in the project, the company went back to relying on the comments of experienced test riders who could describe the experience, in “subjective” terms, of how the motorcycle “felt.”

Judging music-reproduction equipment is no different. No matter how many measurements are gathered about the product’s technical performance, they still don’t tell you how well that product communicates the music. If I had to choose between two unknown DACs as my main source of music for the next five years, I’d rather have ten minutes with each DAC in the listening room than ten hours

with each in the test lab. Today's measurements are crude tools that are inferior to the most powerful test instrument ever devised: the human brain.

A type of vacuum-tubed amplifier called "single-ended triode" (described in [Chapter 5](#)) illustrates the limitations of measurement to describe the musical attributes of audio components. All single-ended-triode amplifiers have technical performance that is, by any standard, laughably bad. They have high distortion, absurdly low output power (often less than ten watts per channel), and can drive very few loudspeakers. Nonetheless, single-ended-triode amplifiers can have a directness of musical communication that must be heard to be believed. Clearly, there are qualities of the audio signal beyond the reach of existing measurements.

Richard Heyser, one of the greatest conceptual thinkers in audio, created a stunningly simple device that revealed the fallacy of relying solely on measurements to judge the quality of audio equipment. The device was a small box with an input jack on one end and an output jack on the other end. When connected to test equipment, the box measured perfectly: it had no distortion, no frequency-response aberrations, and no noise. Indeed, the measurements suggested that all the box contained was a piece of wire connecting input to output.

But when you listened to *music* through the box, the sound was so distorted it was nearly unintelligible. The secret to the Heyser Box is simple: when the box detected the pure, single-frequency sinewaves used in audio testing, a relay inside the box simply connected the input to the output. But when it detected a music signal, with its complex waveform, the relay was repeatedly opened and closed quickly so that the output "chattered." The Heyser Box brilliantly illustrates the danger of relying purely on measurement rather than listening when evaluating audio equipment performance. Moreover, the Heyser Box illustrates the fallacy of testing products in ways that bear no relationship to how those products will be used in the real world.

I must point out, however, that some sophisticated audio designers have developed tests and measurements that are far more advanced than the decades-old parameters routinely used in the audio industry. These advanced measurements can, in some cases, reliably predict specific aspects of audible performance. As their creators are the first to admit, these measurements don't quantify the overall musical experience, nor are they a replacement for critical listening. The measurements are instead tools that help the designer more quickly or more reliably achieve some specifically sought sonic attribute in a product (less treble hardness, for example). Unfortunately, designers consider these

sophisticated techniques trade secrets and don't share them with the audio-engineering community at large.

Given a massive research budget, it may one day be possible to more reliably predict sound quality through measured performance. But that day is a long way off. Until then, we must listen.

Knowing what sounds good and what doesn't is easy; virtually anyone can tell the difference between excellent and poor sound. But discovering *why* a product is musically satisfying or not, and the ability to recognize and describe subtle differences in sound quality, are learned skills. Like all skills, that of critical listening improves with practice: The more you listen, the better a listener you'll become. As your ear improves, you'll be able to distinguish smaller and smaller differences in reproduced sound quality—and be able to describe *how* two presentations are different, and why one is better.

This chapter defines the language of critical listening, describes what to listen for, and outlines the procedures for setting up valid listening comparisons. It will either get you started in critical listening, or help you become a more highly skilled listener.

Audiophile Values

A general discussion of audiophile values is important in understanding the next sections of this chapter. Here are some broad statements about what distinguishes good from superlative sound quality, and what audiophiles value in reproduced music.

Good sound is only a means to the end of musical satisfaction; it is not the end itself. If a neighbor or colleague invites you over to hear his hi-fi system, you can tell immediately whether he's a music lover, or a "hi-fi buff" more interested in sound than in music. If he plays the music very loud, then turns it down after 30 seconds to seek your opinion (approval), he's probably not a music lover. If, however, he sits you down, asks what kind of music *you* like, plays it at a reasonable volume, and says or does nothing for the next 20 minutes while you both listen, it's likely that this person holds audiophile values or simply cares a lot about music.

In the first example, the acquaintance tried to impress you with *sound*. In the second case, your friend also wanted to impress you with his system, but by its ability to express the *music*, not shake the walls. This is the fundamental difference between "hi-fi enthusiasts" and music lovers.

(You can use the same test to immediately tell what kind of hi-fi store you're in. If anyone begins playing the sounds of trains, sonic booms, Shuttle launches, or jet takeoffs, run for cover.)

I've noticed an unusual trend when playing my system for friends and acquaintances not involved in audio. I sit them down in the "sweet spot" (the seating position that provides the best sound) and put on some music I think they'll like. Rather than sitting there for the whole piece or song, they tend to jump up immediately and tell me how good it sounds. They've apparently been conditioned by the kind of razzle-dazzle demonstrations in some hi-fi stores or friends' houses. When *you* begin listening—at someone's house or a dealer's showroom—don't feel the need to express an opinion about the sound. Sit and listen attentively with eyes closed, letting the *music*—*not* the sound—tell you how good the system is.

When listening in a group, don't be swayed by others' opinions. If they're skilled listeners, try to understand what they're talking about. Listen to their descriptions and compare their impressions with your own. In fact, this is the best way to recognize the specific sonic characteristics we'll talk about later in this

chapter. But don't go along with what everyone else says. If you don't hear a difference between two digital cables, for example, don't be afraid to say so. In addition, you should be completely truthful when asked for an opinion about a system. If the sound is bad, say so.

All audio components affect the signal passing through them. Some products add artifacts (distortion) such as a *grainy treble* or a *lumpy bass*. Others subtract parts of the signal—for example, a loudspeaker that doesn't go very low in the bass. (Listening terms are defined later in this chapter.) A fundamental audiophile value holds that sins of commission (adding something to the music) are far worse than sins of omission (removing something from the music). If parts of the music are missing, the ear/brain system subconsciously fills in what isn't there; you can still enjoy listening. But if the playback system adds an artificial character to the sound, you are constantly reminded that you're hearing a reproduction and not the real thing.

Let's illustrate this sins-of-commission/omission dichotomy with two loudspeakers. The first loudspeaker—a three-way system with a 15" woofer in a very large cabinet—sells for a moderate price in a mass-market appliance store; it plays loudly and develops lots of bass. The second loudspeaker sells for about the same price, but is a small two-way system with a 6" woofer. It doesn't play nearly as loudly, and produces much less bass. While you need a refrigerator dolly to move the first loudspeaker, you can almost hold the second loudspeaker in your outstretched hand.

The behemoth loudspeaker has some problems: The bass is boomy, thick, and overwhelming. All the bass notes seem to have the same pitch. The very prominent treble is coarse and grainy, and the midrange has a big peak of excess energy that makes singers sound as if they have colds.

The small loudspeaker has no such problems. The treble is smooth and clear, and the midrange is pure and open. It has, however, very little bass by comparison, won't play very loudly, and doesn't produce a physical sensation of sound hitting your body.

The first loudspeaker commits sins of commission, *adding* unnatural artifacts to the sound. The bass peaks that make it sound boomy, the grain overlaying the treble, and the midrange colorations are all additive distortions.

The second loudspeaker's faults, however, are of omission. It *removes* certain elements of the music—low bass and the ability to play loudly—but leaves the remainder of the music intact. It doesn't add grain to the treble, thickness to the bass, or colorations to the midrange.

There's no doubt that the second speaker will be more musically satisfying. The first speaker's additive distortions are not only much more musically objectionable, they also constantly remind you that you are listening to artificially reproduced music. The second speaker's flaws are of a nature that allows you to forget that you're listening to speakers. In the reproduction of music, addition is far worse than subtraction.

Another audiophile value holds that even small differences in the quality of the musical presentation are important. Because music matters to us, we get excited by *any* improvement in sound quality. Moreover, there isn't a linear relationship between the magnitude of a sonic difference and its musical significance. A quality difference can be sonically small but musically large.

While reviewing a revelatory new state-of-the-art digital-to-analog converter, I listened to a piece of music I'd heard hundreds of times before. The piece, performed by a five-member group, had vocals and very long instrumental breaks. During the instrumental breaks, the vocalist played percussion instruments. Through lesser-quality DACs, the percussion had always been just another sound fused into the music's tapestry; I'd never heard it as a separate instrument played by the vocalist. The group seemed to become a four-piece ensemble when the vocalist wasn't singing; I never heard the percussion as separate from the rest of the music.

The new DAC was particularly good at resolving individual instruments and presenting them not as just more sounds homogenized into the overall musical fabric, but as distinct entities. Consequently, when the instrumental break came, I heard the percussion as a separate, more prominent instrument. In my mind's eye, and for the first time, the vocalist never left—she remained “on stage,” playing the percussion instruments. By just this “small” change in the presentation, the band went from being a quartet to a quintet during the instrumental breaks. The “objective” difference in the electrical signal must have been minuscule; the subjective musical consequences were profound.

This is why small differences in the musical presentation are important—if you care deeply enough about music and about how well it is reproduced. This example highlights the inability of measurements alone to characterize audio equipment performance. Measurements on the DAC in question indicated no technical attributes that would have contributed to my perceptions. More fundamentally, how can a number representing some aspect of the DAC's technical performance begin to describe the musical significance of the change I heard?

Much of music's expression and meaning can be found in such minutiae of detail, subtlety, and nuance. When such subtlety is conveyed by the playback system, you feel a vastly deeper communication with the musicians. Their intent and expression are more vivid, allowing you to more deeply appreciate their artistry. For example, if you compare two performances of Max Reger's Sonata in D Major for solo violin—one competent, the other superlative—you could say that, on an objective basis, they were virtually identical. Both performers played the same notes at about the same tempo. The difference in expression is in the nuances—the inspired subtleties of rubato, tempo, emphasis, articulation, and dynamics that bring the performance to life and convey the piece's musical meaning and intent. This example is analogous to the difference between mediocre and superb music-playback systems, and why small differences in sound quality can matter so much. High-end audio is about reproducing these nuances so that you can come one step closer to the musical expression.

The sad but universal truth about audio equipment is that, any time you put a signal into an audio component, it *never* comes out better at the other end. You therefore want to keep the signal path as simple as possible, to remove any unnecessary electronics from between you and the music. This is why inserting equalizers and other such “enhancers” into the signal chain is usually a bad idea—the less done to the signal, the better. The advent of digital technology, however, has made possible some beneficial signal processing. (An example is digital room correction, described in [Chapter 14](#).)

Pitfalls of Becoming a Critical Listener

There are dangers inherent in developing critical listening skills. The first is an inability to distinguish between critical listening and listening for pleasure. Once started on the path of critiquing sound quality, it's all too easy to forget that the reason you're involved in audio is because you love music, and to start thinking that every time you hear music, you must have an opinion about what's right and what's wrong with the sound. This is the surest path toward a condition humorously known as *Audiophilia nervosa*. Symptoms include constantly changing equipment, playing only one track of a CD or LP at a time instead of the whole record, changing cables for certain music, refusing to listen to great music if it happens to be poorly recorded, and in general "listening to the hardware" instead of to the music.

But high-end audio is about making the hardware disappear. When listening for pleasure—which should be the vast majority of your listening time—forget about the system. Forget about critical listening. Shift into critical-listening mode only when you need to make a diagnostic judgment about the sound quality, or just for practice to become a better listener. Draw the line between critical listening and listening for pleasure—and know when to cross it and when *not* to cross it.

There is also the related danger that your standards of sound quality will rise to such a height that you can't enjoy music unless it's "perfectly" reproduced—in other words, to the point that you can't enjoy music, period. Although it's not very high-quality reproduction, I get a great deal of pleasure from my car stereo—don't let being an audiophile interfere with *your* enjoyment of music, anytime, anywhere. When you can't control the sound quality, lower your expectations.

Sonic Descriptions and their Meanings

The biggest problem in critical listening is finding words to express our perceptions. We hear things in reproduced music that are difficult to identify and put into words. A listening vocabulary is essential not only to conveying to others what we hear, but also to recognizing and understanding our own perceptions. If you can attach a descriptive name to a perception, you can more easily recognize that perception when you experience it again.

A perfect example of the bond between words and perceptions is found in Michael Polanyi's *Personal Knowledge*. His description of the radiology student's experience is identical to that of the audiophile becoming a skilled critical listener:

Think of a medical student attending a course in the X-ray diagnosis of pulmonary diseases. He watches in a darkened room shadowy traces on a fluorescent screen placed against a patient's chest, and hears the radiologist commenting to his assistants, in technical language, on the significant features of these shadows. At first the student is completely puzzled. For he can see in the X-ray picture of a chest only the shadows of the heart and the ribs, with a few spidery blotches between them. The experts seem to be romancing about figments of their imagination; he can see nothing that they are talking about. Then as he goes on listening for a few weeks, looking carefully at ever new pictures of different cases, a tentative understanding will dawn on him; he will gradually forget about the ribs and begin to see the lungs. And eventually, if he perseveres intelligently, a rich panorama of significant details will be revealed to him; of physiological variations and pathological changes, of scars, of chronic infections and signs of acute disease. He has entered a new world. He still sees only a fraction of what the experts can see, but the pictures are definitely making sense now and so do most of the comments made on them. He is about to grasp what he is being taught; it has clicked. Thus, at the moment when he has learned the language of pulmonary radiology, the student will also have learned to understand pulmonary radiograms. The two can only happen together. Both halves of the problem set to us by an unintelligible text, referring to an unintelligible subject, jointly guide our efforts to solve them, and they are solved eventually together by

discovering a conception which comprises a joint understanding of both the words and the things.

Just as the student learns to understand pulmonary radiograms at the moment he learns the language of pulmonary radiology, the audiophile learns to identify specific sonic characteristics when he learns the language of critical listening. The bond between words and ideas is inextricable. That's why listening with more skilled listeners is the best way to learn to identify differences yourself.

By describing in detail the specific sonic characteristics of how electronic components change the sound of music passing through them, I hope to attune you to recognizing those same characteristics when you listen. After reading this next section, listen to two products for yourself and try to hear what I'm describing. It can be any two products—if you have a portable music player, hook it up to your system and compare it to the digital source in your home. The important thing is to start listening analytically. If you don't hear the sonic differences immediately, keep listening. The more you listen, the more sensitive you'll become to those differences.

Each listener hears a little differently. What I listen for during critical listening may be different from another listener's sonic checklist. Moreover, the different values placed on different aspects of the sonic presentation are matters of taste. I notice this firsthand when I occasionally spend time listening critically in my listening room with visiting manufacturers and designers of high-end equipment—many of them highly skilled listeners. While we share many commonalities in determining what sounds good, there is a wide range of perception about what aspects of the presentation are most important.

For example, some listeners value correct reproduction of timbre above all else. A product may have outstanding bass, spectacular soundstaging, and clean treble, but if the timbre is synthetic or hard, those other good qualities won't matter to that listener. Another listener may find soundstaging of utmost importance. Finally, a layer of treble grain that is hardly noticed by one listener may be unacceptable to another. We all hear differently, and value different things in music reproduction. Not only that, but a single listener's musical perceptions and acuity vary with the time of day, as well as his mood and state of relaxation.

There's also evidence that people perceive soundstage information (the impression of music occupying a physical space) in two different ways. When you sit between two loudspeakers reproducing the same signal, your ears each send roughly the same signal to your brain. You thus interpret the sound as coming

from a point exactly between the loudspeakers. If the sound suddenly becomes slightly louder in the left loudspeaker, your brain concludes that the sound source is slightly left of center. This method of localization is called “intensity stereo”: the information your brain uses to localize the sound is the difference in volume, or intensity, between your left and right ears.

A second method of localizing sounds and forming a soundstage in your head uses the phase (timing) information in a recording. Consider our two loudspeakers reproducing the same signal, which you hear as existing between the two speakers. If the left-channel signal is delayed slightly (a few thousandths of a second), you hear the image shift to the right. Although both loudspeakers are producing the same volume of signal, the image appears to come from the loudspeaker that produces the sound first. The two signals are said to have a different phase (timing) relationship to each other.

Some listeners are much more attuned to phase information than others. In recordings made in acoustic spaces (as opposed to recording studios), the phase information is enormously complex. As sounds are reflected by the walls of the acoustic space, they create an almost infinitely dense pattern of very slight delays. It is this temporal (timing) information that cues your brain as to the depth and focus of the images, and the overall size and scale of the concert hall. If a hi-fi system doesn't accurately resolve this temporal information, the soundstage is foreshortened and image placement is vague—to some listeners. Other listeners may not hear the difference between a system that resolves these spatial cues and one that does not. That isn't to say that one way of listening is better than another, only that it appears that two separate mechanisms exist.

The recording that, in my experience, is most revealing of a system's ability to correctly preserve phase information is of John Rutter's *Requiem*, on the Reference Recordings label. When this disc is played back through a first-rate system, one's ability to hear the placement of individual images, along with the sense of the hall's size and characteristics, is stunning. Even a slight degradation of this recording's phase information reduces these qualities. (See [Appendix A](#), “Sound and Hearing,” for more on human hearing.)

The following description of what I listen for is a reflection of my own tastes, mixed with what I've learned from other skilled listeners. You may have your own sonic priorities, but the important thing is that we all hear and describe the same characteristics with the same language. The descriptive terms presented here can be found in my and other writers' product reviews for *The Absolute Sound*. There is general agreement about what these terms mean, but shades of meaning vary among individual listeners.

You should also know that recordings made with audiophile techniques are more revealing of some aspects of reproduced sound than recordings made for mass consumption. For example, a recording of classical music made in a concert hall with very few microphones, a simple signal path, and high-quality recording equipment will likely reveal more about a component's soundstaging performance than a pop recording made in a studio. Similarly, most mass-market recordings have little low bass and almost no dynamic range, so that they sound "good" on a 4" car-stereo speaker. For these reasons, some of the sonic terms described in this chapter apply much more to audiophile-quality recordings than to mass-market ones.

When I listen critically to an audio component to characterize its sonic and musical performance and form judgments about that product, I first follow the setup procedures described later in this chapter. They are vital to forming accurate impressions of products, but secondary to the following discussion.

The sonic terms and characteristics defined here apply to audio components in general. Advice on what to listen for when choosing specific components is included in [Chapters 4–13](#). Note that we often describe differences in terms of differences in frequency response (e.g., "too much treble"). The component may have a flat (accurate) measured treble response, but the distortions it introduces give the impression of too much treble.

It's also useful to understand the broad terms that describe the audio frequency band. The range of human hearing spans ten octaves, from about 16Hz (cycles per second) to 20,000Hz, or 20 kilohertz (20kHz), and can be divided into the specific regions described below. Note that these divisions are somewhat arbitrary; you can't say specifically that the lower treble begins at 2000Hz and not 2500Hz, for example. The table nonetheless provides a rough guideline for understanding the relationship between frequency ranges and their descriptive names.

Frequency Range in Hz

Lower Limit	Upper Limit	Description
16	40	Deep Bass
40	100	Midbass
100	250	Upper Bass
250	500	Lower Midrange
500	1000	(Middle) Midrange
1000	2000	Upper Midrange
2000	3500	Lower Treble
3500	6000	Middle Treble
6000	10,000	Upper Treble
10,000	20,000	Top Octave

This rough guide will help you understand the following terms and definitions. A full characterization of how a product “sounds” will include aspects of each of the following sonic qualities.

Tonal Balance

The first aspect of the musical presentation to listen for is the product's overall *tonal balance*. How well balanced are the bass, midrange, and treble? If it sounds as though there is too much treble, we call the presentation *bright*. The impression of too little treble produces a *dull* or *rolled-off* sound. If the bass overwhelms the rest of the music, we say the presentation is *heavy* or *weighty*. If we hear too little bass, we call the presentation *thin*, *lightweight*, *uptilted*, or *lean*.

A product's tonal balance is a significant—and often overwhelming—aspect of its sonic signature.

Perspective

The term *perspective* describes the apparent distance between the listener and the music. Perspective is largely a function of the recording (particularly the distance between the performers and the microphones), but is also affected by components in the playback system. Some products push the presentation *forward*, toward the listener; others sound more *distant*, or *laid-back*. The forward product presents the music in front of the loudspeakers; the laid-back product makes the music appear slightly behind the loudspeakers. Put another way, the forward product sounds as though the musicians have taken a few steps toward you; the laid-back product gives the impression that the musicians have taken a few steps back.

Another way of describing perspective is by row number in a concert hall. Some products seem to “seat” the listener at the front of the hall—in Row D, say. Others give you the impression that you’re sitting farther back; say, in Row S.

Several other terms describe perspective. *Dry* generally means lacking reverberation and space, but can also apply to a forward perspective. Other watchwords for a forward sound are *immediate*, *incisive*, *vivid*, *palpable*, and *present*. Terms associated with laid-back include *relaxed*, *easygoing*, and *gentle*.

Products with a forward sound produce a greater sense of an instrument’s *presence* before you, but can quickly become fatiguing. Conversely, if the presentation is too laid-back, the music is uninvolved and lacking in immediacy. If the product under evaluation distorts the perspective captured during the recording, I prefer that it err in the direction of being laid-back. When a product is overly immediate, I feel as though the music is coming at me, assaulting my ears. The reaction is to close up, to try to keep the music at arm’s length. The ideal midrange presentation has a sense of *palpable presence*, that feeling of the musicians existing right before you, but without sounding aggressive.

Conversely, a laid-back presentation invites the listener in, pulling her gently forward *into* the music, allowing her the space to explore its subtleties. It’s like the difference between having a conversation with someone who is aggressive, gets in your face, and talks too loudly, compared with someone who stands back, speaking quietly and calmly.

In loudspeakers, perspective is often the result of a peak or dip in the midrange (a peak is too much energy, a dip is too little). In fact, the midrange between 1kHz and 3kHz is called the *presence region* because it provides a sense

of presence and immediacy. The harmonics of the human voice span the presence region; thus, the voice is greatly affected by a product's perspective.

Untrained listeners who can't specifically identify whether a reproduced musical presentation is forward or laid-back will, if the presentation is too forward, feel a tension that usually translates into a desire to turn down the music. Conversely, a laid-back sound can make the music uninvolved rather than riveting.

Note that the terms associated with overall perspective can be used to describe specific aspects of the sound (such as the treble) in addition to describing the overall perspective. If we say the treble is forward, we mean that it is overly prominent, sounding as if it is closer to the listener than the rest of the music.

The Treble

Good treble is essential to high-quality music reproduction. In fact, many otherwise excellent audio products fail to musically satisfy because of poor treble performance.

The treble characteristics we want to avoid are described by the terms *bright*, *tizzy*, *forward*, *aggressive*, *hard*, *brittle*, *edgy*, *dry*, *white*, *bleached*, *wiry*, *metallic*, *sterile*, *analytical*, *screechy*, and *grainy*. Treble problems are pervasive; look how many adjectives we use to describe them.

If a product has too much apparent treble, it overstates sounds that are already rich in high frequencies. Examples are overemphasized cymbals, excessive sibilance (s and *sh* sounds) in vocals, and violins that sound thin. A product with too much apparent treble is called *bright*. Brightness is a prominence in the treble region, primarily between 3kHz and 6kHz. Brightness can be caused by a rising frequency response in loudspeakers, or by poor electronic design. Many digital sources, preamplifiers, amplifiers, and cables that measure as having a flat (accurate) frequency response nevertheless add prominence to the treble.

Tizzy describes too much upper treble (6kHz–10kHz), characterized as a *whitening* of the treble. Tizzy cymbals have an emphasis on the upper harmonics, the sizzle and air that rides over the main cymbal sound. Tizziness gives cymbals more of an *ssssss* than a *sssshhhh* sound. *Forward*, if applied to treble, is very similar to *bright*; both describe too much treble. A forward treble, however, also tends to be dry, lacking space and air around it.

Many of the terms listed above have virtually identical meanings. *Hard*, *brittle*, and *metallic* all describe an unpleasant treble characteristic that reminds one of metal being struck. In fact, the unique harmonic structure created from the impact of metal on metal is very similar to the distortion introduced by a power amplifier when it is asked to play louder than it is capable of playing.

I find the sound of the alto saxophone to be a good gauge of hard, brittle, and metallic treble, particularly lower treble. If reproduced incorrectly, sax can take on a thin, reedy, very unpleasant tone. The antithesis of this sound is *rich*, *warm*, and *full*. When the sax's upper harmonics are reproduced with a metallic character, the whole instrument's sound collapses. Interestingly, the sound of the saxophone has the most complex harmonic structure of any instrument. It's no wonder that it is so revealing of treble problems. *White* and *bleached* have meanings very similar to *bright*, but I associate them more with a thinness in the

treble, often caused by a lack of energy (or what sounds like a lack of energy) in the upper midrange. With no supporting harmonic structure beneath it, the treble becomes threadbare and thin, much like an overexposed photograph. Cymbals should have a gong-like low-frequency component with a sheen over it. If cymbals sound like bursts of white noise (the sound you hear between radio stations), what you're probably hearing is a *white* or *bleached* sound.

A particularly annoying treble characteristic is *graininess*. Treble grain is a coarseness overlaying treble textures. I notice it most on solo violin, massed violins, flute, and female voice. On flute, treble grain is recognizable as a rough or fuzzy sound that seems to ride on top of the flute's dynamic envelope. (That is, the grain follows the flute's volume.) Grain makes violins sound as though they're being played with hacksaw blades rather than bows—a gross exaggeration, but one that conveys the idea of the coarse texture added by grain.

Treble grain can be of any texture, from very fine to coarse and rough. (Think of the difference between 400-grit and 80-grit sandpaper.) The more coarse the grain, the more objectionable it is. The preceding discussion of grain applies in even larger measure to midrange textures, which I'll discuss later.

Treble problems can foster the interesting perception that the treble isn't integrated into the music's harmonic tapestry, but is riding on top of it—the top end seems somehow separate from the music. When this happens, we are aware of the treble as a distinct entity, not as just another aspect of the music. The treble should sound like an extension of the upper midrange rather than separate from it. If the treble calls attention to itself, be suspicious.

The most common sources of these problems are, in rough order of descending magnitude: tweeters in loudspeakers, overly reflective listening rooms, digital source components, preamplifiers, power amplifiers, cables, and noise on the AC powerline.

So far, I've discussed only problems that emphasize treble. However, some products tend to make the treble softer and less prominent than in live music. This characteristic is often designed into the product, either to compensate for treble flaws in other components in the system, or to make the product sound more palatable. Deliberately softening the treble is a designer's shortcut; if he can't get the treble right, he just makes it less offensive by softening it.

The following terms, listed in order of increasing magnitude, describe good treble performance: *smooth*, *sweet*, *soft*, *silky*, *gentle*, *liquid*, and *lush*. When the treble becomes overly smooth, we say it is *romantic*, *rolled-off*, or *syrupy*. A treble described as "smooth, sweet, and silky" is being complimented; "rolled-off and

syrupey” suggests that the component goes too far in treble smoothness, and is therefore *colored*.

A rolled-off and syrupey treble may be blessed relief after hearing bright, hard, and grainy treble, but it isn’t musically satisfying in the long run. Such a presentation tends to become *bland, uninvolving, slow, thick, closed-in, and lacking detail*. All these terms describe the effects of a treble presentation that errs too far on the side of smoothness. If the treble is too soft, the sound will lack *life, air, openness, extension*, and a sense of *space*. The music will sound closed-in rather than big and open.

Top-octave air describes a sense of almost unlimited treble extension, which fosters the impression of “hearing the air” in which the music exists. A slight treble rolloff (a reduction in the amount of treble) in loudspeakers, for example, can diminish top-octave air. Loss of top-octave air is also associated with an opaque soundstage.

The best treble presentation is one that sounds most like real music. It should have lots of energy—cymbals can, after all, sound quite aggressive in real life—yet not have a synthetic, grainy, or dry character. We don’t hear these characteristics in live music; we shouldn’t hear them in reproduced music. More important, the treble should sound like an integral part of the music, not a detached noise riding on top of it. If a component has a colored treble presentation, however, it is less musically objectionable if it errs on the side of smoothness rather than of brightness.

The Midrange

J. Gordon Holt, founder of *Stereophile* magazine and the father of subjective audio-equipment evaluation, once wrote, “If the midrange isn’t right, nothing else matters.”

The midrange is important for several reasons. First, most of the musical energy is in the midrange, particularly the important lower harmonics of most instruments. Not only does this region contain most of the musical energy, but the human ear is much more sensitive to midrange and lower treble than to bass and upper treble. Specifically, the ear is most sensitive to sounds between about 1kHz and 3kHz, and to small changes in both low volume and frequency response within this band. The ear’s threshold of hearing—i.e., the softest sound we can hear—is dramatically lower in the midband than at the frequency extremes. We’ve developed this additional midband acuity probably because the energy of most of the sounds we heard every day for thousands of years—the human voice, rustling leaves, the sounds of other animals—is concentrated in the midrange.

Midrange colorations can be extremely annoying. Loudspeakers with peaks and dips in the mids sound very unnatural; the midrange is absolutely the worst place for loudspeaker imperfections. Confining our discussion to loudspeakers for the moment, midrange colorations overlay the music with a common characteristic that emphasizes certain sounds. The male speaking voice is particularly revealing of midrange anomalies, which are often described by comparisons with vowel sounds. A particular coloration may impart an *aaww* sound; a coloration lower in frequency may emphasize *ooohhh* sounds; a higher-pitched coloration may sound like *eeee*; another coloration might sound *hooty*.

Some midrange colorations can be likened to the sound of someone speaking through cupped hands. Try reading this sentence while cupping your hands around your mouth. Open and close your hands while listening to how the sound of your voice changes. That’s the kind of midrange coloration we sometimes hear from loudspeakers—particularly mass-market ones.

In short, if recordings of male speaking voice sound monotonous, tiring, and resonant, it’s probably the result of peaks and dips in the loudspeaker’s frequency response. (These colorations are most apparent on male voice when listening to just one loudspeaker.)

Terms that describe poor midrange performance include *peaky*, *colored*, *chesty*, *boxy*, *nasal*, *congested*, *honky*, and *thick*. *Chesty* describes a lower-midrange coloration that makes vocalists sound as though they have colds. *Boxy* refers to the impression that the sound is coming out of a small, enclosed space instead of existing in open space. *Nasal* is usually associated with an excess of energy that spans a narrow frequency range, producing a sound similar to talking with your nose pinched. *Honky* is similar to nasal, but higher in frequency and spanning a wider frequency range.

Loudspeaker design has progressed so much in the past ten years that horrible midrange colorations are largely a thing of the past (at least in loudspeakers with high-end aspirations). Still, colorations persist—particularly in inexpensive loudspeakers—though they tend to be more subtle.

One type of lower-midrange coloration that still afflicts even moderately priced high-end loudspeakers is caused by vibration of the loudspeaker's cabinet. The loudspeaker will resonate at certain frequencies, launching acoustic energy from the cabinet when those frequencies are excited by certain notes in the music. Cabinet resonances are heard as certain notes “sticking out” or “jumping forward.” This problem is clearly audible on solo piano; ascending or descending left-hand lines reveal that certain notes possess a character clearly different from the rest of the instrument's notes. These problem notes jump out because they have more energy; some of the sound is being produced by the loudspeaker cabinet, and not just by the loudspeaker's drive units (the cones that move back and forth to create sound). The same thing happens with musical instruments; the musician's expression *wolf tone* describes the same phenomenon. This problem also afflicts the upper bass. In [Chapter 6](#), I talk in much greater detail about loudspeaker cabinet vibration.

As described previously under “Perspective,” too much midrange energy can make the sound seem forward and “in your face.” A broad dip in the midrange response (too little midrange energy over a wide frequency span) can give an impression of greater distance between you and the sound.

When choosing loudspeakers, be especially attuned to the midrange colorations described. What is a very minor—even barely noticeable—problem heard during a brief audition can turn into a major irritant over extended listening.

The preceding descriptions primarily apply to midrange problems introduced by loudspeakers. Expanding the discussion to include electronics (preamps and

power amps) and source components (LP playback or a digital source) introduces different aspects of midrange performance that we should be aware of.

An important factor in midrange performance is how instrumental textures are reproduced. Texture is the physical impression of an instrument's sound—its fabric rather than its tone. The closest musical term for texture is timbre, defined by *Merriam-Webster's Collegiate Dictionary, Eleventh Edition* as “the quality given to a sound by its overtones; the quality of tone distinctive of a particular singing voice or musical instrument.” Sonic artifacts added by electronics often affect instrumental and vocal textures.

The term *grainy*, introduced in the description of treble problems, also applies to the midrange. In fact, midrange grain can be more objectionable than treble grain. Midrange grain is characterized by a coarseness of instrumental and vocal textures; the instrument's texture is granular rather than smooth.

Midrange textures can also sound hard and brittle. Hard textures are apparent with massed voices; a choir sounds *glassy*, *shiny*, and *synthetic*. This problem gets worse as the choir's volume increases. At low levels you may not hear these problems. But as the sound of the choir swells, the sound becomes hard and irritating. Piano is also very revealing of hard midrange textures, the higher notes sounding annoyingly *brittle*. When the midrange lacks these unpleasant artifacts, we say the textures are *liquid*, *smooth*, *sweet*, *velvety*, and *lush*.

A related midrange problem is *stridency*. I think of stridency as a combination of thinness (lack of warmth), hardness, and forwardness, all in the midrange. Strident vocals emphasize mouth noises and overlay the texture with a whitish grain. A saxophone can thus sound thin and reedy, though in a different way from that described as treble grain. A strident sax is hard, more forward in the mids, and granular in texture. Stridency can be caused by an apparent thinness in the lower midrange that makes the upper midrange too prominent. A product described as “strident” has been severely criticized.

Many of these midrange and treble problems are grouped together under the term *harshness*.

Other midrange characteristics affect such areas as clarity, transparency, and detail. These are discussed later, under more specific categories.

The Bass

Bass performance is the most misunderstood aspect of reproduced sound, among the general public and hi-fi buffs alike. The popular belief is that the more bass, the better. This is reflected in ads for “subwoofers” that promise “earthshaking bass” and the ability to “rattle pant legs and stun small animals.” The ultimate expression of this perversity is boom trucks that have absurd amounts of extraordinarily bad bass reproduction.

But we want to know how the product reproduces *music*, not earthquakes. What matters to the music lover isn't *quantity* of bass, but the *quality* of that bass. We don't just want the physical feeling that bass provides; we want to hear subtlety and nuance. We want to hear precise pitch, lack of coloration, and the sharp attack of plucked acoustic bass. We want to hear every note and nuance in fast, intricate bass playing, not a muddled roar. If Ray Brown, Stanley Clarke, John Patitucci, Dave Holland, or Eddie Gomez is working out, we want to hear *exactly* what he's doing. In fact, if the bass is *poorly* reproduced, we'd rather not hear much bass at all.

Correct bass reproduction is essential to satisfying musical reproduction. Low frequencies constitute music's tonal foundation and rhythmic anchor. Unfortunately, bass is difficult to reproduce, whether by source components, power amplifiers, or—especially—loudspeakers and rooms.

Perhaps the most prevalent bass problem is lack of *pitch definition* or *articulation*. These two terms describe the ability to hear bass as individual notes, each having an attack, a decay, and a specific pitch (these terms are explained in [Appendix A](#)). You should hear the texture of the bass, whether it's the sonorous resonance of a bowed double bass or the unique character of a Fender Precision. Low frequencies contain a surprising amount of detail when reproduced correctly.

When the bass is reproduced without pitch definition and articulation, the low end degenerates into a dull roar underlying the music. You hear low-frequency content, but it isn't musically related to what's going on above it. You don't hear precise notes, but a blur of sound—the dynamic envelopes of individual instruments are completely lost. In music in which the bass plays an important rhythmic role—rock, electric blues, and some jazz—the bass guitar and kick drum seem to lag behind the rest of the music, putting a drag on the rhythm. Moreover, the kick drum's dynamic envelope (what gives it the sense of sudden impact) is buried in the bass guitar's sound, obscuring its musical contribution.

These conditions are made worse by the common mid-fi affliction of too much bass.

Terms descriptive of this kind of bass include *muddy, thick, boomy, bloated, tubby, soft, fat, congested, loose, and slow*. Terms that describe excellent bass reproduction include *taut, quick, clean, articulate, agile, tight, and precise*. Good bass has been likened to a trampoline stretched taut; poor bass is a trampoline hanging slackly.

The amount of bass in the musical presentation is very important; if you hear too much, the music is overwhelmed. Excessive bass is a constant reminder that you're listening to reproduced music. This overabundance of bass is described as *heavy*.

If you hear too little bass, the presentation is *thin, lean, threadbare, or overdamped*. An overly lean sound robs music of its rhythm and drive—the full, purring sound of bass guitar is missing, the depth and majesty of double bass or cello are gone, and the orchestra loses its sense of power. Thin bass makes a double bass sound like a cello, a cello like a viola. The rhythmically satisfying weight and impact of bass drum are reduced to shadows of their former power. Instruments' harmonics are emphasized in relation to the fundamentals, giving the impression of well-worn cloth that's lost its supporting structure. A thin or lean sound lacks *warmth* and *body*. As described earlier in this chapter, in the discussion of audio sins of commission and omission, an overly lean bass is preferable to a fat, boomy bass.

Two terms related to what I've just described about the quantity of bass are *extension* and *depth*. Extension is how deep the bass goes—not the bass and upper bass described by *lean* or *weighty*, but the very bottom end of the audioband. This is the realm of kick drum and pipe organ. All but the very best systems *roll off* (reduce in volume) these lowermost frequencies. Fortunately, deep extension isn't a prerequisite to high-quality music reproduction. If the system has good bass down to about 35Hz, you don't feel that much is missing. Pipe-organ enthusiasts will want deeper extension, but correct reproduction of the bottom octave can be very expensive.

Just as colorations caused by peaks and dips in frequency response can make the midrange unnatural, bass can be colored by peaks and dips. Bass colorations create a monotonous, droning characteristic that quickly becomes tiring. In the most extreme example, *one-note bass*, the bass seems to have only one pitch. This impression is created by a large peak in the system's frequency response at a specific frequency. This pitch is then reproduced more loudly than other pitches.

One-note bass is also described as being *thumpy*. Ironically, this undesirable condition is *maximized* in boom trucks. The playback system is tuned to put out all its energy at one frequency, for maximum physical impact. The drivers of those vehicles don't seem to care that they're losing the wealth of musical information conveyed by the bass.

Many of the terms used to describe midrange colorations actually apply to the upper bass. *Chesty*, *thick*, and *congested* are all useful in describing colored bass reproduction. Bass lacking these colorations is called *smooth* or *clean*.

Much of music's dynamic power—the ability to convey wide differences between loud and soft—is contained in the bass. Though I'll discuss dynamics later in this section, bass dynamics bear special discussion—they are that important to satisfying music reproduction.

A system or component that has excellent bass dynamics will provide a sense of sudden impact and explosive power. Bass drum will jump out of the music with startling power. The dynamic envelope of acoustic or electric bass is accurately conveyed, allowing the music full rhythmic expression. We call these components *punchy*, and use the terms *impact* and *slam* to describe good bass dynamics.

A related aspect is *speed*, although, as applied to bass, “speed” is somewhat of a misnomer. Low frequencies inherently have slower attacks than higher frequencies, making the term technically incorrect. But the *musical* difference between “slow” and “fast” bass is profound. A product with fast, tight, punchy bass produces a much greater rhythmic involvement with the music. (This is examined in more detail later.)

Although reproducing the sudden attack of a bass drum is vital, equally important is a system's ability to reproduce a fast decay; i.e., how a note ends. The bass note shouldn't continue after a drum whack has stopped. Many loudspeakers store energy in their mechanical structures and radiate that energy slightly after the note itself. When this happens, the bass has *overhang*, a condition that makes kick drum, for example, sound *bloated* and *slow*. Music in which the drummer used double bass drums is particularly revealing of bass overhang. If the two drums merge into a single sound, overhang is probably to blame. You should hear the attack and decay of each drum as distinct entities. Components that don't adequately convey the sudden dynamic impact of low-frequency instruments rob music of its power and rhythmic drive.

Soundstaging

Soundstaging is the apparent physical size of the musical presentation. When you close your eyes in front of a good playback system, you can “see” the instrumentalists and singers before you, often existing within an acoustic space such as that of a concert hall. The sound-stage has the physical properties of *width* and *depth*, producing a sense of great size and space in the listening room. Soundstaging overlaps with *imaging*, or the way instruments appear as objects hanging in three-dimensional space within the recorded acoustic. As mentioned previously in this chapter, a large and well-defined soundstage is most often heard when playing audiophile-grade recordings made in a real acoustic space such as a concert hall or church.

The most obvious descriptions of the soundstage are its physical dimensions of its width and depth. You hear the musical presentation as existing beyond the left and right loudspeaker boundaries, and extending farther away from you than the wall behind the speakers.

Of all the ways the reproduction of music can be astounding, soundstaging is without question the most miraculous. Think about it: The two loudspeakers are driven by two-dimensional electrical signals that are nothing more than voltages that vary over time. From those two voltages, a huge, three-dimensional panorama unfolds before you. You don’t hear the music as a flat canvas with individual instruments fused together; you hear the first violinist to the left front of the soundstage, the oboe farther back and toward the center, the brass behind the basses on the right, and the tambourine at the very rear, behind all the other instruments. The sound is made up of *individual objects* existing within a space, just as you would hear at a live performance. Moreover, you hear the oboe’s timbre coming from the oboe’s position, the violin’s timbre coming from the violin’s position, and the hall reverberation surrounding the instruments. The listening room vanishes, replaced by the vast space of the concert hall—all from two voltages.

A soundstage is created in the brain by the time and amplitude differences encoded in the two audio channels. When you hear instrumental images toward the rear right of the soundstage, the ear/brain is synthesizing those aural images by processing the slightly different information in the two signals arriving at your ears. Visual perception works the same way: there is no depth information present on your retinas; your brain extrapolates the appearance of depth from the differences between the two flat images.

Audio components vary greatly in their abilities to present these spatial aspects of recordings. Some products shrink soundstage width and shorten the impression of depth. Others reveal the glory of a fully developed soundstage. I find good soundstage performance crucial to satisfying musical reproduction. Unfortunately, many products destroy or degrade the subtle cues that provide soundstaging.

Terms descriptive of poor soundstage width are *narrow* and *constricted*—the music, squeezed together between the speakers, does not envelop the listener. A soundstage lacking depth is called *flat*, *shallow*, or *foreshortened*. Ideally, the soundstage should maintain its width over its entire depth. A soundstage that narrows toward the rear robs the music of its size and space.

The illusion of soundstage depth is aided by resolution of low-level spatial cues such as hall reflections and reverberation. In particular, the reverberation decay after a loud climax followed by a rest helps define the acoustic space. The loud signal is like a flash of light in a dark room: the space is momentarily illuminated, allowing you to see its dimensions and characteristics.

To produce a realistic impression of a real instrument in a real acoustic space, the reverberation and hall sound must be distinct from the image itself. Better audio components place the image within, rather than superimposed on, the recorded acoustic. Poor-quality components don't resolve these spatial cues; they shorten soundstage depth, truncate reverberation decay, and fuse the reverberation into the instrumental images. When this happens, an audio system's ability to transport you to the original acoustic space is diminished.

Now that we've covered space and depth, let's discuss how the instrumental images appear within this space. Images should occupy a specific spatial position in the soundstage. The sound of the bassoon, for example, should appear to emanate from a specific point in space, not as a diffuse and borderless image. The same could be said for guitar, piano, sax, or any other instrument in any kind of music. The lead vocal should appear as a tight, compact, definable point in space exactly between the loudspeakers. Some products, particularly large loudspeakers, distort image size by making every instrument seem larger than life—a classical guitar suddenly sounds 10' wide. A playback system should reveal somewhat correct image size, from a 60'-wide symphony orchestra to a solo violin. I say "somewhat" because it is impossible to re-create the correct spatial perspectives of such widely divergent sound sources through two loudspeakers spaced about 8' apart. Although image size and placement are characteristics inherent in the recording, they are dramatically affected by components in the playback system.

Terms that describe a clearly defined soundstage are *focused, tight, delineated,* and *sharp*. *Image specificity* also describes tight image focus and pinpoint spatial accuracy. A poorly defined soundstage is described as *homogenized, blurred, confused, congested, thick,* and *lacking focus*.

A related issue is soundstage *layering*. This is the ability of a sound system to resolve front-to-back cues that present **some** images toward the soundstage front and at varying distances toward the soundstage rear. The greater the number of layers of depth gradation, the better. Poor components produce the impression of just a few layers of depth: you hear perhaps three or four discrete levels of distance within the soundstage. The best components produce a sense of distance along a continuum: very fine gradations of depth are clearly resolved. Again, the absence of these qualities constitutes subconscious cues that what you're hearing is artificial. When this important spatial information is revealed, however, you can more easily forget that you're not hearing the "real thing."

Bloom—the impression that individual instrumental images are surrounded by a halo of air—is often associated with soundstaging. Although image outlines may be clearly delineated, a soundstage with bloom has an additional sense of diffused air around the image. It is as though the instrument has a little space around it in which it can "breathe." Bloom gives the soundstage a more natural, open, and relaxed feeling.

The term *action*, coined by a colleague at *The Absolute Sound*, Jonathan Valin, describes the sense of bloom expanding outward into space from the instrument, and the way in which the instrument's dynamic envelope grows. Action is bloom with a dynamic component. Think of how a trumpet sounds in real life, and the way the sound seems to expand in space as the instrument plays more loudly. Some audio components express this characteristic better than others; those with a good sense of action sound more realistic, vibrant, and alive than those that lack action.

A product's soundstaging performance should be evaluated with a wide range of musical signals. Some products may throw a superb soundstage at low levels, only to have that stage collapse when the volume increases during musical climaxes. Listen for changes in the spatial perspective with signal level.

Some products produce a crystal-clear, *see-through* soundstage that allows the listener to hear all the way to the back of the hall. Such a *transparent* soundstage has a lifelike immediacy that makes every detail clearly audible. Conversely, an *opaque* soundstage is *thick* or *murky*, with less of an illusion of "seeing" into space. *Veiling* is often used to describe a lack of transparency.

Soundstaging is toward the top of my list of sonic priorities (after tonal balance and lack of grain). The ability to present the music as a collection of individual images surrounded by space, as well as one big image, is very important to the creation of musical realism. A large contributing factor to musical involvement is the impression that the music exists independently of the loudspeakers—that is, when we hear music as existing in space rather than as attached to the position of the speakers. A soundstage with space, depth, focus, layering, bloom, and transparency is nothing short of spectacular.

To evaluate soundstaging and hear the characteristics I've described, you must use recordings that contain these spatial cues. Studio recordings made with multiple microphones and overdubs rarely reveal the soundstaging characteristics described. Recordings made in real acoustic spaces with stereo microphone techniques and a pure signal path are essential to hearing all aspects of soundstaging. In short, soundstaging is provided by a combination of the recording and the playback system. If soundstaging cues aren't present in the recording, you'll never know how well or how poorly the component or system under evaluation reveals them. Most audiophile recordings are made with purist techniques (usually two microphones) that naturally capture the spatial information present during the original musical event.

A term related to soundstaging is *envelopment*, a word that comes from the field of home theater and multichannel sound. Envelopment describes the feeling of immersion in a soundfield created by front and side- or rear-positioned surround speakers. Well-done surround sound can wrap around the listener, creating an almost seamless continuity between the front and rear sounds. (Envelopment is described in more detail in [Chapters 11 and 12](#).)

Finally, superb soundstaging is relatively fragile. You need to sit directly between the speakers, and every component in the playback chain must be of high quality. Soundstaging is easily destroyed by low-quality components, a bad listening room, or poor loudspeaker placement. This isn't to say you have to spend a fortune to get good soundstaging; many very-low-cost products do it well, but it is more of a challenge to find such bargains. [Chapters 14 and 15](#) contain many practical suggestions for achieving the best soundstaging possible from your system.

Dynamics

The *dynamic range* of an audio system isn't how loudly it will play, but rather the *difference* in level between the softest and loudest sounds that the system can reproduce. It is often technically specified as the difference between the component's noise level and its maximum output level. A symphony orchestra has a dynamic range of about 100 decibels, or dB (decibels are defined in [Appendix A](#)); a typical rock recording's dynamic range is about 10dB. In other words, comparatively speaking, the rock band is *always* loud; it has little dynamic range.

Dynamics are a very important part of music reproduction. They propel the music forward and involve the listener. Much of music's expression is conveyed by dynamic contrast, from *pp* (pianissimo) to *fff* (molto fortissimo).

There are two distinct kinds of dynamics. *Macrodynamics* refers to the sound's overall sense of slam, impact, and power—bass-drum whacks and orchestral crescendos, for examples. If the system has poor macrodynamics, we say the sound is *compressed* or *squashed*.

Microdynamics occur on a smaller scale. They don't produce a sense of impact, but are essential to providing realistic dynamic reproduction. Microdynamics describe the fine dynamic structure in music, from the attack of a triangle or other small percussion instruments in the back of the soundstage to the suddenness of a plucked string on an acoustic guitar. Neither sound is very loud in level, but both have dynamic structures that require agility and speed from the playback system.

Products with good dynamics—macro and micro—make the music come alive, allowing a vibrancy and life to emerge. Shifts in dynamics are an important vehicle of musical expression; the more you hear the musicians' intent, the greater the musical communication between performers and listener. Some otherwise excellent components fail to convey the broad range of dynamic contrast.

These characteristics are associated with *transient response*, a system's ability to quickly respond to an input signal. A transient is a short-lived sound, such as that made by percussion instruments. Transient response describes an audio system's ability to faithfully reproduce the quickness of transient signals. For example, a drum being struck produces a waveform with a very steep attack (the way the sound begins) and a fast decay (the way a sound stops). If any component in the playback system can't respond as quickly as the waveform changes, a

distortion of the music's dynamic envelope occurs, and the steepness is slowed. Audio components described as *quick* or *fast* reproduce the suddenness of transient signals. A related concept is *transient coherence*: the impression that the system's transient response is the same across all frequencies. When this happens, bass transients seem lined up in time with midrange and treble transients. Many loudspeakers lack this transient coherence, sounding slower in the bass than in the rest of the audioband.

But just because a component or system can reproduce loud and soft levels doesn't necessarily mean it has good dynamics. We're looking for more than a wide dynamic range. The system must be capable of expressing fine gradations of dynamics, not just loud and soft. As the music changes in level (which, except in most rock recordings, it's doing most of the time), you should hear loudness changes along a smooth continuum, not as abrupt jumps in levels.

Another aspect of dynamics we seek is the ability to play loudly without the sound becoming *congested*. Many products—particularly digital sources and loudspeakers—sound thick during musical climaxes. As the music gets louder, the sound becomes harder, timbres are obscured, and the soundstage degenerates into a confused mess. The word *congeal* aptly describes the homogenization of instrumental images within the soundstage. The nice sense of space and image placement heard at moderate volumes often collapses during loud passages. This produces a sense of *strain* on volume peaks that greatly detracts from your enjoyment. When a component doesn't exhibit these problems, we say it has *effortless* dynamics.

Detail

Detail refers to the small or low-level components of the sound. The fine inner structure of an instrument's timbre is one kind of detail. The term is also associated with transient sounds (those with a sudden attack) at any level, such as those made by percussion instruments. A playback system with good resolution of detail will infuse music with that sense that there is simply more music happening.

Assembling a good-sounding music system or choosing between two components can often be a trade-off between smoothness and the resolution of detail. Many audio components hype detail, giving transient signals an *etched* character. Etch is an unpleasant hardness of timbre on transients that emphasizes their prominence. Sure, you can hear all the information, but the sound becomes too *aggressive, analytical, and fatiguing*: low-level information is brought up and thrust at you, and you feel a sense of relief when the music is turned down or off—not a good sign.

Components that err in the opposite direction don't have this etched and analytical quality, but neither do they resolve all the musical information in the recording. These components are described as overly *smooth*, or having *low resolution*. They tend to make music bland by removing parts of the signal needed for realistic reproduction. These kinds of components don't rivet your attention on the music; they are *uninvolving* and dull. You aren't offended by the sound, as you are with an analytical system, but something is missing that you need for musical satisfaction.

Lack of detail is often an unintended consequence of removing treble brightness or *grain*. The design aspects that tame the aggressive treble also obliterate essential musical information. Such components are also said to have low resolution. Some listeners prefer such low resolution to a more natural sound. In fact, some products are actually designed to gloss over detail to make the products sound more "musical."

It is a rare product indeed that presents a full measure of musical detail without sounding etched. The best products will reveal all the low-level cues that make music interesting and riveting, but not in a way that results in listening fatigue—that sense of tiredness after a long listening session. The music playback system must walk a very fine line between the full resolution of real musical information and sounding analytical.

Pace, Rhythm, and Timing

The British audio reviewer Martin Colloms deserves the credit for identifying and defining important aspects of music reproduction he calls *pace*, *rhythm*, and *timing*. These terms refer to a system's ability to involve the listener in a physical way in the music's forward flow and drive. Pace is that quality that makes your body want to move with the beat, your foot to tap, and your head to bob—the feeling of being propelled forward on the music's rhythm. The sound of good pace and rhythm from an audio system can produce a sense of physical exhilaration.

Though the music's tempo doesn't change in an objective sense from component to component (one power amplifier to another, for example), the subjective impression of a timing difference can be profound. Some products put a drag on the rhythm, making the tempo seem slower or *sluggish*. Others are *upbeat*, conveying the music's rhythmic tautness and drive. In addition to apparently slowing the tempo, components with poor pace give the impression that the band is less tight and musically “up.” The music has less vitality and energy, and the band is less locked-in to the groove.

Colloms wrote this in his revelatory article on this overlooked aspect of music reproduction, published in the November 1992 issue of *Stereophile*.

While good rhythm is a key aspect of both live and reproduced music-making, it is not easy to analyze. It's as if the act of focusing on the details of a performance blinds one to the parameter in question. The subjective awareness of rhythm is a continuous event, registered at the whole-body level, and recognized in a state of conscious but relaxed awareness. Once you've learned that reproduced sound can impart that vital sense of music-making as an event, that the impression of an upbeat, involving drive can be reproduced again and again, you can't help but pursue this quality throughout your listening experience.

Any random energy, such as in a power amplifier's power supply or a loudspeaker's cabinet—even a loosening of the bolts that secure the drivers to the baffles—can degrade pace and timing. Rhythm and pace are more important in rock, jazz, blues, and pop, than in classical music. Much of the feeling of propulsion comes from the bass drum and bass guitar working together.

Many listeners have heard examples of good pace or the lack of it without recognizing it as such. Instead, they may have felt an indefinable involvement in the music if the system had good timing, or a general apathy or boredom with the music if the system didn't convey pace and rhythm.

How do you listen for pace and rhythm? Forget about it; if you find yourself wanting to dance, the component probably has it.

Coherence

Coherence describes the impression that the music is integrated into a satisfying whole, and is not merely a collection of bass, midrange, and treble. The music's harmonic tapestry is integrally woven, not pieced together like a patchwork quilt.

Coherence also applies to the system's or component's dynamic performance, specifically the impression that all the transient edges are lined up with each other. This provides a sense of wholeness to the sound that we react to favorably. Coherence is more a feeling of musical "naturalness" than any specific characteristic.

A related term, coined by Harry Pearson, founder of *The Absolute Sound*, is *continuousness*, the impression that the music is harmonically whole across the audioband.

Musicality

Finally, we get to the most important aspect of a system's sound—*musicality*. Unlike the previous characteristics, musicality isn't any specific quality that you can listen for, but the overall musical satisfaction the system provides. Your sensitivity to musicality is destroyed when you focus on a certain aspect of the presentation; i.e., when you listen critically. Instead, musicality is the *gestalt*, the whole of your reaction to the reproduced sound. We also use the term *involvement* to describe this oneness with the music. A sure indication that a component or a system has musicality is when you sit down for an analytical listening session and, minutes later, find yourself immersed in the music and abandoning the critical listening session. This has happened many times to me as a reviewer, and is a good measure of the product's fundamental "rightness." Ultimately, musicality—not dissecting the sound—is what high-end audio is all about.

The following essay grapples with the conflict between analytical listening and musicality. It first appeared as an editorial in the November 1992 issue of *Stereophile*, and is reprinted with permission.

Between the Ears

by Robert Harley

"Audiophiles constantly seek ways to improve the experience of hearing reproduced music. Preamps are upgraded, digital processors are compared, turntables are tweaked, loudspeaker cables are auditioned, dealers are visited, and, yes, magazines are read—all in the quest to get just a little closer to the music.

"These pursuits have one thing in common: they are all attempts by physical means to enjoy music more. But there's another way of achieving that goal that is far more effective than any tweak, better than any component upgrade, and more fulfilling even than having *carte blanche* in the world's finest high-end store. And it's free.

"I'm talking about what goes on *between* our ears when listening, not what's impinging on them. The ability—or lack of it—to clear the mind of distractions and let the music speak to us has a huge influence on how much we enjoy music. Have you ever wondered why, on the same system and recordings, there is a vast range of involvement in the music? The only variable is our state of mind.

“Because audiophiles care about sound quality, we are often more susceptible to allowing interfering thoughts to get in the music’s way. These thoughts are usually concerned with aspects of the *sound*’s characteristics. Does the soundstage lack depth? Does the bass have enough extension? Is the treble grainy? How does my system compare to those described in magazines?

“Unfortunately, this mode of thinking is perpetuated by high-end audio magazines. The descriptions of a product’s sound—its specific performance attributes—are what make it into print, not the musical and emotional satisfaction to which the product contributes. The latter is ineffable: words cannot express the bond between listener and music that some products facilitate more than others. Consequently, we are left only with descriptions of specific characteristics that can leave the impression that being an audiophile is about dissection and critical commentary, and not about more closely connecting with the music’s meaning.

“A few months after I became an audio reviewer (and a much more critical listener), I underwent a kind of crisis; I found myself no longer enjoying music the way I once did. Listening became a chore, an occupational necessity, rather than the deeply moving experience that made me choose a career in audio. My dilemma was precipitated by the mistaken impression that whenever I heard music I had to have an opinion about the quality of its reproduction. Music became secondary to the sound. Music was merely an assemblage of parts, something to be dismantled and studied, not something that spoke to me emotionally. I found, however, that only after I’d turned in the month’s reviews did the party begin. I stopped being a critic and was a music lover again. I played favorite records, *not* those that would tell me what the equipment was doing. The music as a whole, not the sound as an assemblage of artifacts, once again became the object of my attention. It was as though an enormous burden had been lifted. I made up for lost time in those few days before the inevitable return to reviewing as the next month’s products were set up for critical auditioning.

“Here I was with a custom-built listening room, racks full of the world’s finest audio reproduction equipment, and a job in which I spent a good deal of my time listening, yet most of the time I wasn’t enjoying the music. My cheap car stereo gave more musical satisfaction. Something was dreadfully wrong.

“This experience precipitated a catharsis that forced me to reexamine what music listening—among other things—was all about. I decided to forget about the sound for the vast majority of the listening time and let the music tell me which components were better than others. Critical listening and analytical reasoning became secondary to musical enjoyment. I began listening to music I *liked*, rather

than diagnostic recordings that would tell me about the product's specific characteristics. The impulse to disassemble, listen to the *sound*, and constantly form judgments gradually disappeared. The result is that I now enjoy music more than at any other time of my life. Better sound *does* result in more music, but paradoxically, *only when the sound is forgotten*. Moreover, this conversion increased my listening and reviewing skills: I now have a much better feeling for which products will produce long-term musical satisfaction.

“Nevertheless, analytical listening using particularly revealing recordings must remain a vital part of the reviewing process. Similarly, conveying the specific sonic characteristics of products under review is essential—potential purchasers should know what the product sounds like and decide if it's what they're looking for. But this type of diagnostic analysis no longer dominates—it has become merely one aspect of reviewing. Moreover, I came to realize that the information gained by analytical reasoning is a lower form of knowledge, and that the *feeling* about a product's quality—its ability to convey the music—is a higher form of knowledge.

“These experiences point to a deeper problem of how rationality and the impulse toward dissection as the vehicle to understanding pervades Western thought. Traditional rationality views a whole as a collection of parts. The need to dissect, classify, and assign hierarchical structure is the very foundation of rationality. Our Western upbringing makes it seem so natural that any entity is merely an assemblage of component parts. Why should reproduced music be any different? Consequently, we hear in reproduced music treble, bass, midrange, soundstage, detail, and air. But how often at a live concert do you dismantle the sound the same way you do when listening to a hi-fi? I don't know about you, but I never experience live music in terms of tonal balance, depth, lack of grain, or other characteristics we assign to reproduced sound.

“Consider two approaches to knowing a flower. Traditional rationality would pick it, dissect it, classify its parts, and attempt to understand and document its mechanisms. Another way to know the flower would be to just sit and look at it, appreciating its beauty, discovering subtleties in shape, color, smell, texture, and just absorbing its essence. The first method produces one kind of knowing—an important one, to be sure—but it isn't a complete knowledge.

“Moreover, doesn't the first method destroy something in the process—namely, the flower? After the dissection is finished we find ourselves left with nothing. This is why I no longer enjoyed music; the analytical dismantling consumed the very thing I was trying to know. By allowing intellectual thought to intrude, the oneness between music and listener is destroyed.

“The great accomplishments of rationality—making reproduced music possible in the first place, for example—produces a hubris that tempts one to overlook the virtues of nonrational experience. Rationality can point to its achievements as ‘proof’ of its superiority. The kind of knowing produced by nonrational experience can point to no physical manifestations of its value. Instead, the worth is completely internal and unknowable by anyone who hasn’t experienced it firsthand.

“But it isn’t just listening analytically to the sound that interferes with musical enjoyment. *Any* intellectual activity diminishes the experience. It is the quality of the listening experience that matters—the here and now—not an assignment of value as measured against some standard promulgated by hi-fi magazines, comparisons with other systems, or anything else that distracts from the music.

“There’s nothing wrong with being dissatisfied with your system; that’s what fuels improvement. And a better system does produce a deeper musical experience. But you shouldn’t enjoy music any less if your system falls short of a friend’s system, or if you don’t have the latest hot product, or if the sound doesn’t equal that described in magazines. These things are artificial inventions *superimposed on* reality, not the reality itself. What really matters is the experience—immediate, here-and-now existence. All other thoughts or intellectual abstractions diminish our potential for a oneness with the music.

“In this way of thinking (or, more precisely, *not* thinking), musical fulfillment is no longer dependent on external means. Yes, the technology that brings music into our homes is physical and external, but satisfaction cannot be experienced purely through the equipment, no matter how good it is. Instead, getting closer to the music is a special interaction between music and listener, the equipment merely serving as intermediary. Without this state of mind, no level of equipment will provide us with what we seek. My temporary lack of musical enjoyment, despite having a custom-built room filled with top-notch products, affirmed this truism.

“By all means, continue critiquing and upgrading your system—it can be a path to more musical enjoyment. But when a favorite record goes on the turntable and the lights dim, forget about cables, preamps, tweaks, and magazines. At that moment, the only thing that *really* matters is the music.”

Notes on Learning Descriptive Terms

By practicing critical listening, and adopting the language used here and in product reviews in high-end audio magazines, you will be much better able to assess the quality of audio equipment and make better purchasing decisions. By associating the descriptive terms with your own listening impressions, you can more tangibly characterize how good products and systems are, and *why* they are good or bad.

I recommend that you write down your listening impressions during or immediately after the listening session. This will not only solidify your impressions, but force you to find words to describe what you heard. Often we remember more strongly the memory of the impression, rather than the impression itself. I'm often asked how I can remember how a specific component sounds when I haven't heard that product for a year or more. Because I record my sonic experiences when writing product reviews for *The Absolute Sound*, I have formed a capsule impression of that product's basic character. I don't remember every detail, of course, but the mental image may be something like "slightly etched treble, very transparent and spacious soundstage, lean bass, somewhat lacking in dynamics, analytical rather than smooth." I also associate a general value judgment with that component: would I want to listen to it for pleasure over a long period of time?

By writing down these impressions, you'll not only have a written record of your experience, but also a more tangible mental impression, one that you can use when making later comparisons with other components.

Critical-Listening Setup Procedures

Now that we've explored the lexicon of critical listening, let's look at the procedures required to reach valid listening judgments. These procedures are controls on the process that help ensure that our listening impressions are valid. It's easy to be fooled if you don't take certain precautions.

The following procedures are the result of my experience evaluating products professionally, and reporting those evaluations to a sizable audience. In my profession, the need to be right about products is paramount. Not only will readers buy or not buy products based on my opinions, but the fortunes of high-end manufacturers are greatly affected by reviews, positive or negative. Accurately describing a product's sound, and reaching correct judgments about that product, are huge responsibilities—to readers, manufacturers, and the truth. The techniques described are essential to reaching the right conclusions. I wouldn't even consider forming concrete opinions about products without these safeguards. Neither should you.

By using these controlled procedures, you can stop second-guessing what you hear and concentrate on the real musical differences between components. Removing the variables that may skew the listening impressions will allow you to have greater confidence that the differences you're hearing are real musical differences between products, not the results of improper setup. If you follow these procedures and hear differences, you can have confidence in your listening impressions.

As you go into the listening session, turn your mind into a blank slate. Forget about the brand names of the products you're about to audition, their reputation for certain sonic characteristics, the products' prices, and whom you may offend if you like product A better than product B. Be completely receptive to whatever your ears tell you. If your impressions match your preconceived ideas about what the product sounds like, suspect bias. But if you find that your listening impressions contradict your prejudices, you've developed the ability to listen without bias. Removing bias is essential to reaching accurate judgments.

The first rule of critical listening is that only one variable should be changed at one time. If comparing preamps, use the same source components, interconnects, AC power conditioner, power amp, cables, loudspeakers, music, and room. This is obviously impossible when comparing products at different dealers with different playback systems. But changing one variable at a time is an

absolute requirement for arriving at valid opinions. This is why I stress taking the competing products home over the weekend for comparison. Not only can you evaluate them under the same conditions, but you get to hear how they interact with the rest of your system.

Changing only one component at a time is vital; big differences (two different pairs of loudspeakers, for example) will often swamp small differences. You can't separate what one new component is doing from what the second new component is doing. The ultimate expression of this mistake is someone who hears a system for the first time and remarks on how good (or bad) the loudspeaker cables are. It's ridiculous; you can't isolate what the loudspeaker cable is doing within the context of a completely unknown system. An exception to this rule is when evaluating AC power cords. It is often more revealing to change all the AC cords in the system at once rather than one at a time.

The second rule is to match volume levels between components under audition. Level matching ensures that both products produce the same volume at the loudspeakers, and should become an integral part of critical listening sessions. All it takes is an inexpensive voltmeter and a few minutes of time. If you're comparing two digital sources, for example, match their output levels to within, at most, 0.2dB of each other, and preferably to within 0.1dB. You'll be well rewarded for those few minutes spent level-matching with more accurate judgments about the products. (The section at the end of this chapter explains how to match levels.)

Level-matching is crucial; slight level differences between products can lead you to the wrong conclusion. As described in [Appendix A](#), the ear's sensitivity to bass and treble increases disproportionately with volume. That is, you hear more bass and treble when the music is loud. If product A is played louder than product B, product A may sound brighter, more forward, more detailed, more dynamic, and have more bass. If product A has a soft treble, lacks detail, and has a lean bass balance, you may not know it if the levels aren't precisely matched. If you *can* distinguish product A's character with mismatched levels, the degrees to which the treble is soft, the detail lacking, the dynamics missing, and the bass thin are much harder to gauge. With level matching, you'll know precisely how *much* these characteristics are present in the product you're evaluating; you can forget about trying to mentally compensate for different levels and concentrate on characterizing the sonic differences.

With the same playback system and matched levels, you're ready to compare products. The typical method is called A/B, in which you hear product A, then immediately afterward the same music through product B. One problem with

A/B listening is called “The A/A Paradox.” This states that when hearing two identical presentations (A/A), you tend to hear differences between them. The A/A paradox happens because music has meaning; you perceive it and react to it a little differently each time you hear it. Moreover, if the music is unfamiliar, a second hearing allows you to hear more detail. These factors combine to fool you into hearing differences that aren’t there.

Fortunately, overcoming the A/A paradox is easy. When comparing products, listen to A, then B, then A again. The differences you heard between A and B are often solidified after hearing A again. After selecting another piece of music for further comparison, make the order B/A/B. This technique of hearing three presentations (or five, A/B/A/B/A) confirms or refutes first impressions. Usually, a characteristic heard in the first comparison is more apparent in a second comparison.

Selecting the right *source material* (i.e., recordings) is important in fully exploring a product’s character. If the source material doesn’t contain certain cues, you won’t know how well the product reproduces those cues. For example, if you’re comparing two power amplifiers and use only dry studio recordings of pop music, you’ll never know how good the amplifiers are at resolving soundstage depth and other spatial cues. Similarly, if you listen to nothing but chamber music in evaluating loudspeakers, you can’t tell how the loudspeaker conveys bass dynamics, punch, and rhythmic intensity.

Establish a repertoire of familiar music for evaluating components, each recording selected for some sonic characteristic that will tell you how a component behaves. Remember, this is diagnostic listening, not listening for pleasure. A range of source material should include: a full-scale symphonic recording, chamber music, orchestra with choir, solo piano, popular vocal, rock, blues, or other music with electric bass and kick drum, jazz with acoustic bass, and music with cymbals. Some of the jazz and orchestral recordings should be naturally miked; that is, recorded in a way that captures the original spatial aspects of the music and concert hall. Audiophile record labels will often describe how the recordings were made, and even include a photograph of the hall and where the microphones were placed. Get to know these recordings intimately. By hearing the same recording played back on many different systems, you can quickly characterize a component’s sonic signature.

Another factor that can lead you astray in critical listening is *absolute polarity* (discussed in [Appendix A](#)). If one product under comparison inverts absolute polarity and the other one doesn’t, you may be led to the wrong conclusions; or, at the very least, be confused about which product is better. Many owner’s

manuals will state if the product inverts polarity. If you know that product A does not invert polarity while product B does, reverse the positive and negative connections on *both* loudspeakers: the red wire goes to the loudspeaker's black terminal, and the black wire goes to the loudspeaker's red terminal when auditioning product B. If your preamp or DAC has a polarity-inversion switch, you can just throw it instead of swapping loudspeaker leads. Fortunately, very few components invert absolute polarity.

It's also easy to be misled by a colored playback system. Say you have loudspeakers that are bright, forward, and a little rough in the treble, and you're comparing two digital sources to determine which one to buy. Digital source A produces a sound through your system that is bright and forward in the treble. Digital source B renders a smoother treble and is more musical in your system.

Is source B inherently more musical? Not necessarily. It probably has colorations—an overly soft treble—that complement the loudspeaker's bright treble. Two products in the playback chain that have reciprocal flaws will often produce a musical result. And there's no reason not to buy source B if you intend on keeping your loudspeakers. But don't make the mistake of thinking source B is inherently better than source A. In fact, source A is probably more neutral and would be a much better choice with different, smoother loudspeakers. Although not ideal, choosing a product based on how it corrects other faults in your system is valid if the end result is a more musical sound. Just recognize the situation when it occurs.

Single-Presentation Listening—What It’s All About

The type of analytical listening I’ve described will reveal specific sonic attributes of components in relation to each other: DAC A has deeper bass extension, smoother treble, and a deeper soundstage, while DAC B has more inner detail, quicker transients, and a more focused soundstage. This kind of analytical listening will tell you a lot about the products’ strengths and weaknesses, but it isn’t the whole story. What’s missing from this dissection is how each product makes you *feel* about the music. Does your mind wander, thinking about what you’re going to do next? Or are you riveted in the listening seat, playing record after record well into the night?

This ability to express the music in a way that compels you to continue listening is called a product’s *musicality* and *involvement*. Musicality and involvement are rarely revealed in A/B comparisons and can’t be described in terms of specific sonic attributes. Instead, this fundamental characteristic of a product’s ability to provide long-term musical satisfaction is discovered only when listening for pleasure, without level-matching, selecting music for its sonic rather than musical qualities, or analytical thinking. Dissecting the music to characterize specific sonic attributes, as in A/B comparisons, destroys our receptivity to the music’s meaning. Involvement comes only when you consider music as a whole, not as a collection of sonic parts.

This is why all critical listening should include not only A/B/A comparisons, but also what’s called “single-presentation listening” and listening just for pleasure. Single-presentation listening is listening to the product under evaluation over a long time period—days or weeks rather than hours—with an unconscious ear to how the product affects the music. Listening for sheer pleasure involves no thinking about the sound, only about the music. If, after a long session, you feel exhilarated and fulfilled, you can be confident in the product’s ability to convey the music’s expression and meaning. In fact, this is, without question, the most important indicator of quality in an audio component.

There is a general correlation between a product’s sonic performance as revealed during analytical listening and its musicality as discovered during listening for pleasure. Specific sonic flaws—grainy treble, hard textures, flat soundstage—will often distract you from the music in ways that prevent deep involvement. The less the presentation sounds artificial, synthetic, or affronts

your ears, the easier it is to forget the hardware and hear only the music. Indeed, making the hardware disappear is the holy grail of high-end audio.

Your evaluative listening should thus include a mix of analytical A/B/A comparisons with single-presentation listening (with an ear to specific sonic attributes) and pure listening for pleasure.

Critical Listening Summary

I've summarized the setup procedures for conducting critical listening that will provide valid, useful information about the product under evaluation:

- *Use the same playback system for comparisons*
- *Change only one component at a time*
- *Match levels between products under audition*
- *Listen to A/B/A, not A/B*
- *Use very familiar recordings*
- *Select music that reveals specific sonic characteristics; e.g., bass, midrange liquidity, dynamics, treble, and soundstaging*
- *Listen to the product for pleasure over a long period of time*

Do I follow all these guidelines every time I listen critically? No. But I don't reach firm opinions about products unless I do. If you're listening to make a purchase decision, follow these guidelines religiously. Be suspicious of any listening impression received in poor circumstances. And don't reach firm opinions unless all of these guidelines have been followed.

By applying these techniques, you'll be more likely to reach accurate value judgments, and consequently buy products that will provide the greatest long-term musical satisfaction.

Addendum: Level Matching

Level-matching procedures are identical for all electronic components, whether a digital source, preamp, or power amp. You'll need a voltmeter (even a cheap one will work) and a test CD that contains pure tones at various frequencies. First, play music at a comfortable level through the system to set the volume. Then play the track on the test CD of a 1kHz sinewave recorded at a low level. Measure the voltage across the loudspeaker terminals when the tone is playing, and note the voltage. Put a piece of masking tape on the preamp's face plate around the preamp's volume control and mark the position of the volume control.

After you switch components (power amp, digital source, etc.), play the 1kHz sinewave again and measure the voltage across the loudspeaker terminals. Adjust the preamp's volume control until the voltage across the speaker terminals is the same as what you measured for the first component. Mark the volume-control position on the masking tape. By setting the volume control at the first mark for product A and at the second mark for product B, you've got matched levels.

Some preamplifiers with detented (click stop) volume controls but continuous attenuation make level matching easier and more precise by removing the need for tape. Just count the number of steps necessary to achieve the same voltage across the speaker terminals. Setting the volume control between detents is surprisingly accurate and repeatable.

If you're comparing preamps, you don't need masking tape—just find the volume setting on both preamps that produces the same voltage across the loudspeaker terminals when playing the test tone, and leave the volume controls in the same position. Matching levels for phono cartridges and phono preamplifiers is identical, except that a test LP replaces the test CD.

Level-matching isn't required when comparing products that don't affect the volume. These include interconnects, speaker cables, digital interconnects, AC line conditioners, AC power cords, vibration-damping devices, room treatments, and equipment stands.

Ideally, the playback level between products should be matched to within 0.2dB, and preferably to within 0.1dB. To calculate the dB difference between two measured voltages, you need only a calculator with a "log" button. Here's how to do it.

First, find the ratio of the two measured voltages by dividing one by the other. Then find the logarithm of that ratio by pressing the calculator's "log" button with the ratio in the display. Next, multiply the result by 20. The answer is the difference in dB between the two voltages.

For example, if you measured 2.82 volts (V) across the loudspeaker terminals with DAC A, and 2.88V across the speaker terminals with DAC B after level matching, you first divide 2.88 by 2.82, which gives you the ratio of the two voltages (1.0213:1). Next, push the "log" button on the calculator with 1.0213 in the display, and multiply the result, 0.00914, by 20 to get the answer: the levels are matched to within 0.18dB, or close enough for meaningful comparisons.

4 Preamplifiers

Introduction

The preamplifier is the Grand Central Station of your hi-fi system. It receives signals from source components—DACs, music servers, turntables, disc players—and allows you to select which of these to send to the power amplifier for listening. In addition to allowing you to switch between sources, the preamplifier performs many other useful functions, such as amplifying the signal from your phono cartridge (in some preamplifiers), adjusting the balance between channels, and allowing you to set the volume level. The preamplifier is the component you will most often use, touch, and adjust. It also has a large influence on the system's overall sound quality. In an integrated amplifier or receiver, the preamplifier is built into (integrated with), the other components in a single housing.

In addition to providing volume control and letting you select which source component you listen to, the preamp is a *buffer* between your source components and power amplifier. That is, the preamplifier acts as an intermediary, taking in signals from source components and conditioning those signals before sending them on to the power amplifier. Source components can easily drive a preamplifier, with the burden of driving a power amp through long cables falling on the preamp's shoulders. By buffering the signal, the preamp makes life easier for your source components and ensures good technical performance.

So far I've been describing the preamplifier and its function in a classic audio system. But new technologies have imbued some preamplifiers with innovative features and capabilities that greatly expand the preamp's role in music playback.

At a minimum, some of today's preamplifiers include an integral digital-to-analog converter (DAC), obviating the need to buy a separate DAC. The inclusion of a DAC within the preamplifier housing is analogous to the inclusion of the ubiquitous phono stage within preamps designed before the rise of CD.

(Preamps with phono stages are, however, making a comeback, fueled by the LP's spectacular resurgence.)

But integral digital-to-analog conversion is just the start of the preamplifier's metamorphosis into a product category almost entirely new—a component that acts as the hub of an advanced Internet-connected digital-audio system. Let's look at some of these newer preamplifier features.

Streaming music services: Listen to music *streamed* from subscription services such as Tidal, Spotify, Rhapsody, and others.

Internet radio: Tune into just about any radio station in the world.

Network connectivity: Connect the preamplifier to a local *network* and stream music from any device on the network, such as a *network-attached storage* (NAS) drive.

Wireless connectivity: Stream music from a portable device to your home-audio system via the preamplifier's wireless capability. Wireless options may include one or more technologies, including *WiFi*, *Bluetooth*, and *near-field communication* (NFC). (All of these technologies are described in detail in [Chapter 8](#).)

As you can see, new technologies are blurring the lines between traditional product categories. Moreover, these features aren't unique to preamplifiers; they can be found alone or together in a variety of products. Is a component that provides volume control, source selection, and digital-to-analog conversion a preamplifier with a DAC or a DAC with some preamplifier functions? How about a component that offers network connectivity along with volume control and source selection? Is a component that is supplied with a music-management app still a preamplifier?

I consider a component to be a preamplifier if it provides analog inputs, analog source selection, and a volume control, no matter what its other features. A DAC with a volume control and the ability to select between digital sources isn't a preamplifier because it lacks analog inputs and analog source switching. This definition isn't so much an industry standard as my view of each component's role in a music-playback system.

Fig.4-1 shows the front panel of one of these newer breeds of preamplifiers. Notice the *USB input* and lack of traditional buttons and knobs, along with the large front-panel display. The model is configured and controlled via a single knob under software control. The rear panel (**Fig.4-2**) looks very different from that of a traditional preamplifier, with many more and varied inputs and connections. Note also that this product's digital inputs are mounted on removable boards (the vertical modules on the left). This modular design allows you to upgrade the hardware as new digital interface standards are introduced.



Fig.4-1 A new breed of preamplifier offers networking capability, wireless streaming, and integration with streaming-music services. (Courtesy NAD)



Fig.4-2 The rear panel of the model shown in **Fig.4-1** looks very different from that of a traditional preamplifier. (Courtesy NAD)

Keep in mind that preamplifiers that aspire to the absolute state of the art of sound quality are unlikely to include these capabilities. Rather, such products' mission is to do three things—source selection, volume control, act as a buffer between your sources and power amplifier—without compromise. Additional

features can degrade sound quality; the parts budget must be spread out over much more circuitry, connections, and chassis space, not to mention the potential for noise from the digital circuits polluting the analog signal path. Given the same size case, a traditional preamplifier can devote more space to the power supply and analog circuits than can a “do-it-all” product. If you have the budget and rack space, buy a traditional preamplifier and then add the newer music-accessing technologies with additional components. This approach also allows you to keep your preamplifier, and to upgrade the digital components as technology inexorably advances.

Having said that, there are compelling arguments for today’s multifunction preamps: lower cost, a simpler system, less “system clutter,” and greater ease of use. Just be aware that, at the upper end of the price and performance scales, the traditional preamplifier still reigns supreme sonically.

There are many types of preamplifiers, each with different capabilities and functions. Choosing the one best suited to your system requires you to define your needs. Listeners without a turntable, for example, won’t need a preamplifier that amplifies the tiny signals from a phono cartridge. Others will need many inputs to multiple source components, and some will need multichannel capability. Let’s survey the various preamplifiers and define some common preamplifier terms.

Line-Stage Preamplifier: Accepts only line-level (low-level) signals, which include every source component except a turntable. Line stages have become much more popular as listeners increasingly rely on digital sources rather than LPs as their main signal source. If you don’t have a turntable, you need only a line-stage preamplifier.

Phono Preamplifier (also called a phono stage): Takes the very tiny signal from your phono cartridge and amplifies it to line level. It also performs *RIAA equalization* on the signal from the cartridge. RIAA equalization is a bass boost and treble cut that counteract the bass cut and treble boost applied in disc mastering, thus restoring flat response. Some phono preamplifiers offer other equalization curves besides RIAA (more on this later).

A phono stage can be an outboard stand alone unit in its own chassis, or a circuit section within a full-function preamplifier. If you play records, you must have a phono stage, either as a separate component or as part of a full-function preamplifier.

Full-Function Preamplifier: Combines a phono stage with a line-stage preamplifier in one case.

Preamplifier/DAC: A preamplifier (with or without a phono stage) that includes an integral digital-to-analog converter.

Tubed Preamplifier: Uses vacuum tubes to amplify the audio signal.

Solid-State Preamplifier: Uses transistors to amplify the audio signal.

Hybrid Preamplifier: Uses a combination of tubes and transistors.

Audio/Video Controller: A device analogous to a preamplifier that includes video switching, multiple audio channels (typically eight), and surround-sound decoding such as Dolby Digital and DTS. (A/V controllers are described in [Chapter 11](#), “Audio for Home Theater.”)

Multichannel Preamplifier: A preamplifier with multiple audio channels (typically six) for playback of multichannel music. Differs from an A/V controller in that the multichannel preamplifier has no surround decoding, video switching, or other functions for film-soundtrack reproduction.

Digital Preamplifier: A preamplifier that accepts digital input signals (such as from a CD transport) and processes the audio signal in the digital domain. Digital preamplifiers usually include a digital-to-analog converter so that the digital preamplifier can drive a conventional power amplifier.

Passive Level Control: Sometimes erroneously called a passive preamp, the passive level control can replace a line-stage preamplifier in some situations. It is inserted in the signal path between a source component and the power amplifier.

Instead of amplifying the source, a passive level control merely attenuates (reduces) the signal level driving the power amplifier. It doesn't plug into the wall, and cannot amplify a signal as does an active line-stage preamplifier.

No Preamplifier: Some DACs can drive a power amplifier directly with no need for a preamplifier. The DAC must have a variable output level so that you can

adjust the volume. Some DACs also offer multiple digital inputs along with source switching, obviating the need for a preamplifier.

How to Choose a Preamplifier

Once you've decided on a line-stage, a full-function preamp, or separate line and phono stages (the last are generally more expensive), it's time to define your system requirements. The first is the number of inputs you'll need. Count your source components and be sure that the preamplifier has that number of inputs plus one (to accommodate future needs).

Most high-end preamps have few features—and for good reason. First, the less circuitry in the signal path, the purer the signal and the better the sound. Second, the preamp designer can usually put a fixed manufacturing budget into making a preamp that either sounds superb or has lots of features—but not both. Mass-market mid-fi equipment emphasizes vast arrays of features and buttons at the expense of sound quality. Don't be surprised to find very expensive preamps with almost no features; they were designed, first and foremost, for the highest sound quality. Most high-end preamps don't even have tone (bass and treble) controls. Not only do tone controls electrically degrade the signal—and thus the musical performance—but the very idea of changing the signal is antithetical to the values of high-end audio. The signal should be reproduced with the least alteration possible. Tone controls are usually unnecessary in a high-quality system.

Another school of thought, however, holds that a playback system's goal isn't to perfectly reproduce what's on the recording, but to achieve the most enjoyable listening experience possible. If changing the tonal balance with tone controls enhances the pleasure of listening to music, use them. If you're comfortable with the latter philosophy, be aware that tone controls invariably degrade the preamplifier's sound quality. (Some preamps with tone controls do allow you to switch them out of the circuit when they're not being used, a feature called *tone defeat*.)

A similar debate rages over whether or not to include a balance control(s) in a high-end preamp. A balance control lets you adjust the relative levels of the left and right channels. If the recording has slightly more signal in one channel than the other, the center image will appear to shift toward the louder channel, and the sense of soundstage layering may be reduced. A similar problem can occur if the listening room has more absorptive material on one side than the other, pulling the image off-center. A small adjustment of the balance control can correct these problems. Like tone controls, balance controls can slightly degrade a

preamplifier's sonic performance. Some upper-end preamplifiers provide a balance control with no sonic penalty. It isn't unusual to find a \$25,000, state-of-the-art preamplifier with no tone or balance controls, and a \$199 mass-market receiver with both of these features.

If you plan to use your 2-channel preamplifier as part of a multichannel home-theater system, you'll need a feature called *theater bypass* (or similar name). This is an input on the preamplifier that simply routes the signal at the preamplifier's input jacks to the preamplifier's output jacks with no gain or attenuation, and with no ability to adjust the volume. This feature is necessary in a hybrid system in which a 2-channel stereo system is augmented with a multichannel controller, a multichannel power amplifier, and center and surround speakers. (This system configuration is explained in detail in [Chapter 11](#).)

Two other preamplifier features you should know about are *12V trigger* and *RS232 control*. Let's first look at the simpler 12V trigger. When you turn on a preamplifier by pressing the preamp's front-panel power button, 12V DC appears at the 12V trigger output jack. When connected to another component with a 12V trigger input (such as a power amplifier), the preamplifier's 12V trigger signal turns on the other components. Many components can be daisy-chained together in this way (provided that each has a trigger input and output), so that pressing just one power button turns on the entire system. This feature is not only more convenient, it allows you to position some components in cabinets or other inaccessible locations.

RS232 control is considerably more sophisticated. Components with RS232 ports can be completely controlled by a computer, sophisticated aftermarket remote controls, or home-automation systems such as those made by Crestron and AMX. A preamplifier with RS232 can be placed in a different room, yet you'll still be able to control all aspects of it through the home-automation system's touchscreen interface or your mobile device. The automation system sends commands to the preamplifier through the RS232 interface. This scenario is popular in custom installations in which the audio equipment is completely hidden, controlled by a touchscreen.

A preamplifier feature recently making a strong comeback is the integral headphone amplifier. You can immediately identify a preamp with this feature by the $\frac{1}{4}$ " *phone jack* on the front panel. The headphone amplifiers built into preamplifiers vary greatly in quality. Some are included merely for convenience; others are designed with the same care and attention as the preamp's main circuits. If headphone listening is important to you, choose a preamp with a high-

quality headphone amp. If you're really serious about headphone listening, consider a dedicated, stand alone headphone amplifier. (See [Chapter 13](#) for an in-depth look at headphones and headphone amplifiers.)

When making a purchasing decision, you should also consider the preamplifier's look and feel. Because the preamplifier is the component you'll interact with most, ask yourself: Are the controls well laid out? Is the volume control easy to find in the dark? Does the preamp have a mute switch? A mute switch is handy when your listening is interrupted, for protecting the rest of your system when disconnecting cables, or for maintaining the same volume level when comparing two other components, such as digital sources or cables. Another useful feature is an *absolute polarity switch* (Absolute polarity is described in [Appendix A](#).)

Note that these recommendations for choosing the preamp best suited to your system also apply to integrated amplifiers. Think of an integrated amplifier as a preamplifier that also happens to have a power amplifier built in. An integrated amplifier is functionally nearly identical to a preamplifier. The only difference is that, instead of connecting the preamplifier output to a power amplifier, you attach speaker cables to the integrated amplifier. (See [Chapter 5](#) for more information on integrated amplifiers.)

Balanced and Unbalanced Connections

Some preamplifiers have *balanced* inputs, balanced outputs, or both. A balanced signal is carried on a three-pin XLR connector rather than the conventional RCA plug. (Balanced and unbalanced connections are described in [Chapter 10](#).) [Fig.4-4](#) later in this chapter shows the rear panel of a preamplifier with balanced inputs and outputs on XLR jacks, and unbalanced inputs and outputs on the familiar RCA jacks. If you have a balanced source component—usually a CD player or DAC—you should consider choosing a preamplifier with a balanced input. Nearly all source components with balanced outputs also have unbalanced outputs. You can use either output, but you may not get the best sound quality unless you use the balanced output option. (See [Chapter 7](#), “Digital Source Components,” for the reason why.) In addition, not all “balanced” preamplifiers are created equal, as we’ll see later in this chapter.

A preamp with unbalanced inputs and a balanced output can accept unbalanced signals but still drive a power amplifier through a balanced interconnect. If your power amplifier has balanced inputs, getting a balanced-output preamplifier is a good idea. You can listen to the system through both balanced and unbalanced lines and decide which sounds better. Some products sound better through balanced connections; others perform best with unbalanced lines. (A more technical description of balanced preamplifiers appears later in this chapter.)

What to Listen For

The preamplifier has a profound effect on the music system's overall performance. Because each of the source signals must go through the preamp, any colorations or unmusical characteristics it imposes will be constantly overlaid on the music. You can have superb source components, a top-notch power amplifier, and excellent loudspeakers, yet still have mediocre sound if the preamp isn't up to the standard set by the rest of your components. The preamplifier can establish the lowest performance level of your system; careful auditioning and wise product selection are crucial to building the best-sounding playback system for your budget.

A preamplifier's price doesn't always indicate its sonic quality; I know of one \$2000 model that is musically superior to another preamp selling for nearly \$6000. If you do your homework and choose carefully, you'll avoid paying too much for a poor-sounding product.

In addition to the usual listening procedures described in [Chapter 3](#), preamplifiers offer several methods of sonic evaluation not possible with other components. We can therefore more precisely evaluate preamps and choose the best one for the money.

We'll start with the standard listening-evaluation techniques. First, the same musical selection can be played on the same system, alternating between two competing preamps. Be sure to match levels between the two preamps under audition. Listen for the presentation differences described in [Chapter 3](#)—particularly clarity, transparency, lack of grain, low-level detail, soundstaging, and a sense of ease. The better preamplifier will “get out of the way” of the music, imposing less of itself on the sound.

The most common sonic problems in preamplifiers are hardening of timbres, a thickening of the soundstage, and a reduction in low-level detail. Many preamps, particularly inexpensive solid-state models, overlay the midrange and treble with a steely, metallic hardness. These preamps can give the impression of more musical detail, but quickly become fatiguing. The treble becomes drier, more forward, and etched. Cymbals lose their sheen, instead sounding like bursts of white noise. Vocal sibilants (s and *sh* sounds) become objectionably prominent; violins become screechy and thin. The poor-sounding solid-state preamplifier emphasizes the brightness of the strings and diminishes the

resonance of the violin's wooden body. Such preamps also reduce the saturation of tonal color, sounding thin and bleached.

The preamplifier that thickens the sound makes the soundstage more opaque. The transparent quality is gone, replaced by a murkiness that obscures low-level detail and reduces resolution. Instruments and voices no longer hang in a transparent, three-dimensional space. Instead, the sound is thick, confused, congealed, and lacks clarity. Even some expensive models impose these characteristics on the music.

Beware of the preamplifier that tends to make all recordings sound similar in timbre, tonal balance, or spatial presentation. The best preamplifiers let you hear changes in the sizes of concert halls, how closely the microphones were positioned, and can clearly resolve the differences between Steinway and Bosendorfer pianos, among other such subtleties. Preamplifiers also vary greatly in their resolution, or ability to reveal very fine musical detail. The low-resolution preamplifier seems to shave off a layer of musical information at the quiet end of the dynamic spectrum. Compare the preamps under audition to the very best preamp in the store. Listen for the qualities distinguishing the best preamp, and see if those characteristics are in the preamps you're considering. This will not only give you a reference point in selecting a preamplifier for yourself, but will sharpen your listening skills. The more experience you have listening to a variety of products, the better your ability to judge component quality.

A useful way of evaluating preamplifiers—one not possible with other components—is the *bypass test*. This technique compares the preamplifier to no preamplifier, revealing the preamplifier's editorial effect on the music. This most revealing of comparisons leaves the preamp's shortcomings nowhere to hide (Fig.4-3).

To conduct a bypass test, drive a power amplifier at a comfortable volume with a known high-quality preamplifier, or directly from a DAC that has an output level control. Using a voltmeter and a test CD with test tones, set the preamplifier under evaluation for unity gain—the volume position at which the input level is identical to the output level (usually a setting a third to halfway up the volume control's range). Put the preamplifier under evaluation in the signal path between the known preamp and the power amp. Compare the system's sound with and without the preamp in the playback chain. Did the preamp thicken the sound? Did the treble become dry and brittle with the preamp in the system? Was the sense of space, clarity, and transparency replaced by an opaque thickness that homogenized the individual instruments? Did the preamplifier impart a common sonic characteristic to varied recordings? Most important, was

the music less involving? The bypass test provides a quick and accurate assessment of the type and amount of coloration imposed by a preamplifier.




Fig.4-3 The bypass test for evaluating preamplifiers

If you're comparing two preamplifiers, use preamp A to drive the power amplifier; listen, then put preamp B in the signal path at unity gain between preamp A and the power amplifier. Listen for the difference between the sounds with and without preamp B in the signal path. Then reverse the roles of the two preamps, listening for preamp A's effect on the sound. This technique quickly reveals exactly what each preamp does to the music. This method, however, doesn't provide listening conditions of the highest resolution; the colorations of the first preamplifier can obscure qualities in the second preamplifier. It also adds another run of interconnects and terminations to the signal path.

You can also conduct a bypass test by substituting a passive level control for the preamplifier. Although passive level controls can reduce dynamic contrast and soften the bass in some systems, they rarely add grain and etch, or reduce soundstage depth and clarity, as do most active preamplifiers.

The increasing prevalence of digital source components with a variable output level and the ability to directly drive a power amplifier provides a test bed for evaluating preamplifier transparency. This is especially true when listening to files of very high-resolution (176.4kHz/24-bit) decoded by state-of-the-art DACs; the high-res material is extremely revealing of any colorations downstream. In my

recent experience of listening to such sources with and without a preamplifier in the signal path, I've found that even heroic (and expensive) preamplifiers noticeably degrade the sound. Timbres are not quite as pure, the soundstage is slightly less expansive, and the very lowest levels of fine detail are obscured. If you have a DAC with a variable output level, I encourage you to perform the bypass test.

A less analytical method is to borrow the preamp from your dealer for a weekend and just listen to it. How much more exciting and involving is the music compared to using your existing preamp? Does the new preamp reveal musical information you hadn't heard before in familiar recordings? How much does the new preamp compel you to continue playing music? These are the best indicators of the product's ability to provide long-term musical satisfaction. Trust what your favorite music tells you about the preamplifier.

Tubes vs. Transistors

Of the many components that make up a hi-fi system, the preamplifier is the most likely to use vacuum tubes instead of solid-state devices (transistors). This is because preamplifiers handle only low-level signals, making tubes more practical and affordable there than in power amplifiers. Power-amplifier tubes are large, expensive, run hot, and require periodic replacement. Further, one theory of audio design holds that if the system is to include tubes, they are best employed closest to the signal source—such as in the preamplifier.

Moreover, the qualities that have endeared many music lovers to the magic of tubes are much more affordable when used in a preamplifier. Tubed audio components are more expensive to build and maintain than solid-state designs, but that cost differential is much lower in preamplifiers than in power amplifiers. If you want the special qualities of tubes but not their heat, greater expense, and higher maintenance, the tubed preamplifier is the way to go.

Tubes are often claimed to sound sweeter and warmer, portray greater soundstage dimensionality, and to have a more natural treble. Many solid-state preamps tend to make the treble sound dry, brittle, metallic, and etched. The result is steely-sounding strings (particularly violins), unnaturally emphasized vocal sibilants (*s* and *sh* sounds), and cymbals that sound like bursts of high-frequency noise rather than a delicate brass-like shimmer. Because these unpleasant artifacts can be introduced by many components (digital sources, power amplifiers, cables, tweeters), a natural-sounding tubed preamplifier can tend to ameliorate the system's tendency toward these amusical characteristics.

The popularity of tubed electronics has nothing to do with nostalgia and everything to do with sound quality. Tubes have several advantages in their favor as audio amplifying devices. First, the circuitry associated with a vacuum tube—the ancillary parts that make the tube work—is generally much simpler than a circuit using transistors. Second, the distortion tubes produce is very different from the distortion created by solid-state electronics. Harmonic distortion, introduced by all active electronics, adds spurious frequencies to the signal being amplified. For example, if a preamplifier is passing a 1kHz signal, the preamp will generate distortion at 2kHz (second harmonic), 3kHz (third harmonic), 4kHz (fourth harmonic), and so forth. Tubed circuits generally produce lower-order harmonic-distortion components such as second and third harmonic, in contrast with the higher-order distortion (seventh and ninth harmonic) introduced by

transistors. Lower-order distortion is much less musically objectionable than upper-order. Moreover, a large amount of second-harmonic distortion (say, 5%) is harder to hear than a very small amount (say, 0.5%) of seventh-harmonic distortion. Tube proponents point to these facts as evidence of the superiority of tubed preamplifiers.

Be aware, however, that some tubed preamps are intentionally designed to sound very colored. Rather than present the music with the least added effect, tubed preamps often add significant amounts of *euphonic coloration*. This form of distortion can at first be pleasing to the ear, but represents a departure from the original signal. This type of “tubey” coloration is characterized by a soft treble, an overly laid-back and easygoing sound, lack of detail, and a “syrupy” quality. Many audiophiles, in their attempts to avoid the worst characteristics of solid-state, turn to euphonicly colored tubed preamps to make their systems listenable.

It is far better, however, to make sure that each component in the chain is as transparent as possible. If this is achieved, there will be no need for “tubey” preamps. Ideally, the listener shouldn’t be aware that she is listening to tubes; instead, she should be aware only of the music. Just as poor solid-state preamps color the sound by adding grain, treble hardness, and etch, the poor tube preamp will often err in the opposite direction, obscuring detail, adding false “bloom,” and reducing resolution. Be equally aware of both forms of coloration.

It’s a mistake to “fall in love” with either solid-state or tubes for the wrong reasons. The solid-state lover may think he is getting “more detail,” and the tube aficionado may fall for the “lush sweetness.” In the pursuit of a truly musical playback system, both extremes are to be avoided. The overly “sweet” preamplifier may become uninvolved over time because of its low resolution; the “detailed” and “revealing” solid-state model may eventually become unmusical for the fatigue it produces in the listener. Not all tubed and solid-state preamps can, however, be categorized so neatly. Many tubed preamps are extremely transparent and neutral, having very little effect on the sound.

Nor should you buy a preamp *only* because it uses tubes. Certain circuits are better implemented with tubes, others with solid-state devices. There are no magic components or circuit designs that will ensure a product’s musicality. Worthy—and unworthy—products have been made from both tubes and transistors. That’s why some companies design tubed, solid-state, and hybrid tubed/solid-state preamplifiers. The designer picks the best device for the particular application.

Ideally, both tubed and transistor preamplifiers strive for musical perfection, but approach that goal from different directions. The gross colorations I've described are largely things of the past, or relegated to fringe products that appeal to a minority of unsophisticated listeners. Today's reality is that tubed and solid-state designs sound more and more alike every year. The best preamplifier will be a transparent window on the music—no matter what its design.

The best advice: Choose the preamp that has the least effect on the music; you'll get much more musical satisfaction from it for a longer period of time. And remember: The perfect tubed preamp and the perfect solid-state preamp would sound identical.

Tube Life and Replacement Options

The small-signal tubes found in preamplifiers (which generally use from two to eight tubes) need replacing after 1000 to 2000 hours of use. Any time the preamplifier is turned on, the hour meter is ticking. This fact, along with the heat produced by tubes, encourages many owners of tubed preamplifiers to turn their units off when not listening. Although this extends tube life, an electronic component doesn't sound its best until it's been on for a few hours. The best advice is to turn on the preamplifier an hour or so before a listening session, if possible. If you know you'll be listening at night after work, turn on the preamplifier in the morning and it will be at peak performance when you get home. Another alternative is to buy a mechanical AC timer (about \$10) from a hardware store and set it to switch on the preamp two hours before you expect to come home. Solid-state preamplifiers, on the other hand, draw very little current, run cool, and don't have tubes to replace—they can be left on continuously.

Replacement tubes vary in cost from \$10 for an untested generic type to \$100 for a premium-grade tube that has been thoroughly tested and selected for best electrical performance. You may also want to consider paying a premium for matched pairs of tubes that have the same gain (amount of amplification). If the preamplifier's channel balance (similarity of gain between left and right channels) is determined by the gain of separate tubes, it's worth the extra cost for matched pairs. This ensures that one channel won't be louder than the other.

Because tube quality has a big effect on the preamplifier's sound, high-quality tubes are worth the money. New Old Stock (NOS) tubes (tubes that are unused but that haven't been manufactured in a long time) carry the highest prices because of limited supply and high demand. Different brands of tubes equal in technical performance will also sound different, making choosing difficult. Your local dealer is the best guide in selecting the optimal replacement tubes for your preamplifier. Reviews of tubed preamplifiers will sometimes include the reviewer's opinions about different tube brands used in a particular preamplifier. Changing tubes in an amplifier is called "tube rolling."

Fortunately for the tube aficionado, the supply of reliable, high-quality tubes has never been better. In fact, the increasing popularity of tubed audio products is partially due to the recent availability of great tubes from the former Soviet Union. Their technology was heavily reliant on vacuum tubes (which are still used in MIG fighter jets), and their factories produced excellent products. With the end of the Cold War, this large supply of superb, reasonably priced tubes has

fueled the recent growth of the tubed audio industry. The Chinese tube industry has also stepped up its output and quality in recent years.

The Line-Stage Preamp

A typical minimalist line-stage preamplifier is shown in Fig.4-4. This model has four stereo pairs of unbalanced inputs plus four stereo pairs of balanced inputs. A front-panel input selector switch selects which input is routed to the output. A mute switch shuts off the signal at the preamplifier output. (This is useful when changing records or disconnecting and connecting source components—engaging the mute switch prevents a loud thump from being sent through your speakers.) Note that this preamplifier has no balance control for adjusting the relative levels between left and right channels. Two sets of main output jacks are provided, along with a “record output” whose level is fixed regardless of the volume-control setting. The front-panel controls include a switches for mono operation, polarity inversion, and source selection.



Fig.4-4 A line-stage preamp’s front and rear panels. (Courtesy Constellation Audio)

A pair of 12V input/output jacks allow the preamp to be turned on by another component, or to output 12V to turn on other components. This 12V “trigger”

signal is widely used as a way of simplifying powering up or down an entire system. For example, a disc player's 12V trigger output would drive the preamplifier, automatically turning on the preamplifier when the disc player is turned on. The preamplifier's 12V trigger output in turn feeds the power amplifier's 12V trigger input, turning on the power amplifier when the preamplifier is turned on. This 12V trigger system allows you to turn on your entire system with one button push.

The Phono-Stage Preamplifier

The phonograph (phono) stage amplifies the very tiny signal (between a few tens of microvolts and a few millivolts) from the phono cartridge to a line-level (about one volt) signal. This line-level signal can then drive a line-stage preamplifier, just as any other source component would. A phono stage can be an integral part of a full-function preamplifier, an optional board that plugs into some line-stage preamplifiers, or an outboard unit in its own case. Outboard phono stages have no volume controls; they usually feed a line input on a line-stage preamplifier.

In the days before CD, virtually all preamplifiers included integral phono stages. In many of today's preamplifiers, however, a phono stage is an option (usually about \$200–\$800) for those listeners who play LPs. This arrangement reduces the preamp's price for those needing only a line-stage preamplifier. It's also less expensive to buy a preamplifier with an integral phono stage than buying separate line and phono stages. With one case, one power supply, one owner's manual, and one shipping carton, a full-featured preamp is less expensive. These integral phono stages can also offer exceptional performance.

Not all line stages accept a phono board; if you think you'll want to play records in the future, be certain the line-stage preamp you buy will accept a phono board. If the preamp's rear panel has an input marked Phono, a ground lug, and a front-panel input selection position marked Phono, you can be confident that the preamp has a built-in phono stage or will accept a phono board. Note, however, that some line-stage preamps have a line-level input marked Phono but no phono stage. Look on the rear panel for the ground lug—a post for connecting a ground wire between the preamplifier and the turntable. If the preamp has a ground lug, it probably has provisions for accepting phono signals.

Some phono stages are simple boxes with a pair of RCA inputs, a pair of RCA outputs, and a power cord. Others offer multiple inputs, remote control, a front-panel display of gain and cartridge loading (described later in this chapter), and balanced inputs or outputs (Figs. 4-5 and 4-6).



Fig.4-5 A simple, inexpensive phono stage. (Courtesy Pro-Ject and Sumiko)



Fig.4-6 An elaborate phono stage with multiple inputs, variable cartridge loading, front-panel gain adjustment, balanced outputs, and remote control. (Courtesy Aesthetix)

RIAA Equalization

In addition to amplifying the cartridge's tiny output voltage, the phono stage performs RIAA equalization on the signal. RIAA is the Recording Industry Association of America, the body that standardized record and playback equalization characteristics. Specifically, phono-stage RIAA equalization boosts the bass and attenuates the treble during playback. This equalization counteracts the bass cut and treble boost applied to the signal when the record was cut. By combining exactly opposite curves in disc mastering and playback, a flat response is achieved. RIAA equalization curves are shown in Fig.4-7. The dotted trace is the curve applied when the lacquer master is cut; the solid trace is the curve applied in the phono preamplifier. The combination of these complementary curves is flat response.

Attenuating bass and boosting treble when the disc is cut allows more signal to be cut into the record groove and increases playing time. Because bass takes up more room in the groove than does treble, reducing the amount of bass in the groove allows the grooves to be spaced closer together. In addition, the treble boost in disc cutting and subsequent treble cut on playback act as a noise-reduction system: attenuating treble on playback also attenuates record-surface noise. RIAA equalization is why you can't plug a line-level source into a preamplifier's phono input. Similarly, a phono cartridge can't be plugged into a preamplifier's line input.

One aspect of a phono preamp's technical performance is its RIAA accuracy. This measurement or specification indicates how closely the phono stage's RIAA equalization circuit matches the standard curve used in disc mastering. Typical performance is 20Hz–20kHz, ± 0.5 dB. That means the phono preamplifier will exhibit frequency-response variations (some frequencies are emphasized over others) of no more than 1dB (plus or minus half a dB) over the audioband of 20Hz to 20kHz. Because RIAA errors are essentially frequency-response errors, they can greatly affect the phono stage's sound. A positive RIAA error in the treble can make the phono stage sound bright and etched; a negative error in the treble can make the sound dull and lifeless.




Fig.4-7 RIAA equalization curves.

Note that European audio equipment often has a low-frequency rolloff curve built into the phono stage's RIAA circuit. That is, low frequencies are intentionally reduced in volume. This prevents the very low frequencies (below about 30Hz) generated by warped records from getting into the rest of the system, but can also reduce bass extension on playback systems capable of reproducing frequencies down to 20Hz.

In the late 1990s, a few phono stages began appearing that offered alternate phono equalization curves, on the premise that some older LPs might have been cut with an EQ curve other than RIAA. Before the RIAA standard was established in 1954, different record labels used their own equalization curves. If those recordings were played back with today's standard RIAA curve, the result would be a tonal imbalance, such as too much treble. Selectable equalization curves allow the listener to apply the correct curve and hear a flat frequency response. These adjustable preamplifiers usually include a "Columbia" curve and a "Decca" curve, in addition to RIAA.

If you don't listen to records made before 1954, you probably don't need a selection of curves. All record labels converted to RIAA in the mid- to late 1950s; playing an RIAA LP with a curve other than RIAA (which introduces greater treble rolloff) is a deviation from neutrality, but one that flatters the many overly bright records of that era. Some listeners, however, have suggested that proprietary equalization was used at some labels long after 1954, and that playing back those LPs with a Columbia or Decca curve results in more accurate sound.

Phono-Stage Gain

The amount of amplification provided by a phono stage (or any other amplifier) is called its *gain*. Gain is specified either in decibels (dB), or as a number expressing the ratio between the input and output voltages. Phono stages have much more gain than line stages. Where a line stage may have 10 to 20dB of gain, a phono stage typically amplifies the signal by 40 to 75dB.

The amount of phono-stage gain required depends on the type of phono cartridge driving the phono preamplifier. Phono stages are of two varieties, each named for the type of cartridge with which it is designed to work. The first is the moving-magnet phono stage. Moving-magnet phono stages have their gain optimized to work with the relatively high output voltages produced by moving-magnet cartridges. Moving-magnet cartridges have high output voltages as cartridges go, on the order of two to eight millivolts (2–8mV). Consequently, they need less gain; a moving-magnet phono stage's gain is toward the lower end of the range, typically about 35dB.

Moving-coil cartridges have much lower output voltages due to their different method of generating a signal. Moving-coil output levels range from 0.15 to 2.5mV. Consequently, they need more amplification (gain) in the phono preamplifier to reach line level than do signals from moving-magnet cartridge. Moving-coil phono preamplifiers have about 35–68dB of gain. Note that moving-coil output voltages vary greatly with the cartridge design, with some high-output models reaching moving-magnet levels.

Because of this wide variation in cartridge output level, a gain mismatch can occur between the cartridge and phono stage: The phono stage can have either too much or too little gain for a specific cartridge's output voltage. If the phono preamp doesn't have enough gain, the volume control must be turned up very high for sufficient playback levels. This raises the noise floor (heard as a loud background hiss), often to an objectionable level. Conversely, a high-output cartridge can overload a moving-coil phono stage's input circuitry, causing distortion during loud musical passages. This condition is called input overload. A high-output cartridge driving a high-gain phono preamplifier can also make the preamplifier's volume control too insensitive. A moderate listening level may be achieved with the volume control barely cracked open; this makes small volume adjustments difficult. This latter condition may be accompanied by distortion on musical peaks: the phono stage's excessive output level causes the line stage's input circuitry to overload on peaks.

Correctly matching the cartridge's output voltage to the phono-stage gain avoids excessive noise and the possibility of input overload. A moving-coil cartridge specified at 0.18mV output needs about 75dB of gain. A typical moving-magnet output of 3mV should drive a phono stage that has about 50dB of gain. Some phono stages and full-function preamps have internal switches that adjust the gain between moving-magnet and moving-coil levels. The best phono stages allow you to select from many gain settings in fine increments (say, 3dB per step), so that you can precisely match the gain to your particular cartridge. An example of an elaborate system of gain and loading adjustment is shown in Fig.4-8. These "DIP switches" are located on the bottom panel of this phono stage, and provide a wide range of gains and cartridge loadings.

A pre-amplifier is a small, stand alone component that can boost a moving-coil signal up to moving-magnet levels. If you have a moving-magnet phono input and upgrade to a moving-coil cartridge, you can add a pre-amplifier instead of getting a new phono stage. Similarly, a step-up transformer increases a moving-coil's output voltage to a higher level. A transformer can improve the phono system's signal-noise ratio, or allow a moving-coil cartridge to drive a moving-magnet phono input.

High phono-stage gain carries the penalty of increased noise, particularly in all-tube phono stages. Some tubed phono stages use a transistor (usually a Field Effect Transistor, or FET) in the first amplification stage, followed by subsequent tubed stages. This technique greatly increases the product's signal-noise ratio, but at the expense of adding a solid-state gain stage to the signal path. The amount of noise in a tubed phono stage is greatly influenced by the particular tubes in the unit. Replacing noisy tubes can greatly reduce the noise level. Although phono stages in general have poorer signal-noise specifications than other components, very-high-gain phono preamps can be objectionably noisy. All other factors being equal, the greater the gain, the higher the noise. Select a phono preamp with just enough gain for your cartridge.



Fig.4-8 This phono stage offers extensive gain and loading options on its bottom panel. (Courtesy Simaudio)

For the mathematically inclined, here's a formula for precisely calculating the optimal phono-stage gain for a given cartridge output voltage. The formula is $NdB = 20 \log V1/V2$. "NdB" is the number of decibels of gain, "V1" is 1 volt (the desired output voltage from the phono stage), "V2" is your cartridge's output voltage, and "20 log" is 20 times the logarithm (base 10) of the ratio between V1 and V2. We'll use a cartridge with an output voltage of 0.7mV in our example. We first divide 0.0007 into 1 and get the ratio 1428. The logarithm of 1428 (you can use an online log calculator) is 3.155. We then multiply 3.155 by 20 to arrive at the optimal gain of 61dB.

Why is 1V the ideal output voltage from the phono stage? A cardinal rule of audio is that gain followed by attenuation degrades the signal-noise ratio. A typical power amplifier driving a loudspeaker of typical sensitivity produces a normal listening level when driven by a signal of about 1V. If the phono stage output level is 1V, the line stage between the phono stage and the power amplifier operates at unity gain (no gain or attenuation) and thus introduces the minimum amount of noise itself.

Cartridge Loading

Cartridge loading is the resistance and capacitance the phono cartridge is presented with when driving the phono input. Cartridge loading, specified in both impedance and capacitance, has a large effect on how the cartridge sounds, particularly moving-magnet types. Improper loading can cause frequency-response changes and other undesirable conditions. Many preamps allow you to adjust the input impedance and input capacitance to match the phono cartridge by adding resistors and capacitors to the phono stage's input circuit. Cartridge loading is done via a front-panel switch or knob, replacing boards inside the phono stage, adjusting DIP switches on or inside the phono stage, or even replacing parts with a soldering iron. Some phono stages offer variable cartridge loading via remote control from the listening seat. This latter feature is extremely convenient. Note that the tonearm cable's capacitance must be added to the phono preamplifier's input capacitance to find the total capacitance loading the cartridge.

The Digital Preamplifier

Although the “digital preamplifier” never caught on as an audiophile product, I’ve included a description here for the sake of completeness. A digital preamplifier takes in digital signals from a source component such as the USB output from a computer, adjusts the volume and performs other signal processing in the digital domain, then converts the digital signal to analog for output to your power amplifier.

The term “preamplifier” is really a misnomer because the device we’re talking about doesn’t “preamplify” signals. It could just as well be called a digital-to-analog converter that has built-in source switching, volume control, and sometimes tone and balance adjustments. The product is, however, called a “digital preamplifier,” and is built to look and feel like a traditional analog preamp because that’s the product it replaces. Just as we “open files” in our computers and adjust the “leading” of type, the digital product we’re describing is a “preamplifier.”

Because it operates in the digital domain, the digital preamplifier can perform some types of signal processing we’d never attempt in the analog domain, such as dynamic-range compression/expansion, soundstage-width control, and radical equalization. These processes are largely to correct recordings that have sonic deficiencies; most of the time, you leave them switched out of the circuit.

Digital preamplifiers can have a few analog inputs that bypass all the digital circuitry (and forgo the signal processing just described) and function as a traditional preamplifier. Some, however, digitize analog input signals, then convert them back to analog for output to your power amplifier. These A/D and D/A conversions inevitably degrade sound quality.

If you decide that a digital preamplifier is right for you, an important feature is a variable input level on its analog inputs. This is a knob on the front panel accompanied by a meter or LEDs that display the input level, which controls the level feeding the preamplifier’s A/D converters. The ability to control the level feeding the A/D converter is crucial: If the level is too low, you’ll hear noise and distortion; if too high, the signal will distort on loud peaks.

Audio/Video Controllers and Multichannel Preamplifiers

The popularity of home theater, with its multichannel sound, has enticed many makers of high-end preamps to produce models that work in 2-channel stereo systems as well as those incorporating multichannel surround sound. These products are called A/V (for audio/video) preamplifiers or A/V controllers. The latter term is more technically correct, and better describes the product's function. Although A/V controllers are covered in depth in [Chapter 11](#), a brief introduction here is illustrative.

A/V controllers are distinguished by at least four factors: multichannel sound capability, video switching, surround decoding, and bass management. First, an A/V controller will typically have eight audio channels instead of two, to accommodate the 7.1 channels of audio in today's surround-sound formats. Second, the A/V model can switch video as well as audio sources. A/V controllers all perform surround decoding, such as Dolby Pro Logic, Dolby Digital, DTS, Dolby Digital Plus, DTS HD, Dolby TrueHD, and DTS HD Master Audio. Finally, all A/V controllers offer bass management: the ability to filter bass information from some channels and direct that bass to a subwoofer.

A key feature of a high-end A/V controller is the ability to pass analog input signals to the output without converting the signal to digital and then back to analog. This analog bypass feature avoids unnecessary sonic degradation imposed by A/D and D/A conversions.

For those who want multichannel music reproduction without video switching and surround decoding, the multichannel preamplifier is a better choice than an A/V controller. A multichannel preamplifier is simply a conventional preamplifier with six (or eight) channels rather than two. Although it offers no surround decoding, bass management, or other home-theater-related functions, it will likely offer better sound quality than a digital controller, for several reasons. First, it contains no digital circuitry, which creates noise that can get into the analog signal. Second, the designer can focus the entire parts budget on making the product sound good without worrying about including all that other complex circuitry. Very few multichannel preamplifiers are available, since they are only for those who will use their systems exclusively for listening to music.

How a Preamplifier Works

A preamplifier consists of several stages. Each is an electronic circuit that performs a specific function, and can be pictured as a building block of an entire preamp.

The preamplifier's input stage is the first circuit element in the signal path. It acts as a buffer between the preamp's internal circuitry and the components driving it. That is, the input stage presents a high input impedance to source components driving the preamp, and a low output impedance to the next preamplifier stage. Impedance is, broadly speaking, the resistance to current flow. You'd intuitively think that more resistance to electrical current flow (a higher input impedance) in the preamplifier would put greater stress on the source components driving the preamp. The opposite is true: the higher the preamplifier's input impedance, the less current the source component must deliver.

Think of a garden hose connected to a faucet. The water pressure is the voltage from the source component, the hose is the interconnect cable, and your thumb and finger squeezing the hose are equivalent to the preamplifier's input impedance. The harder you squeeze the hose, the greater the resistance (impedance) reducing the amount of water flowing through it. If you let go of the hose altogether (a very low input impedance), the water pressure drops because the water supply is limited. In an audio system, some digital-source components with limited current output will distort (or have soft bass) when asked to drive a low-input-impedance preamplifier. The same digital source components will sound just fine when driving a high-input-impedance preamplifier. The preamplifier's high-impedance input stage is designed to make life easy for the source components driving the preamp.

The next circuit stage provides gain, amplifying the input signal. This is where the preamplifier actually boosts the signal level. (Although the input stage can have some gain.) The last stage is the output stage, the circuit block that serves as a buffer between the gain stage and the power amplifier connected to the preamp. The output stage has a high input impedance, low output impedance, and is designed to drive interconnects and a power-amplifier input. These circuit blocks are shown in [Fig.4-9](#). Each of these stages can be made from transistors, vacuum tubes, integrated circuits (operational amplifiers, or op-amps), or a combination of these devices.

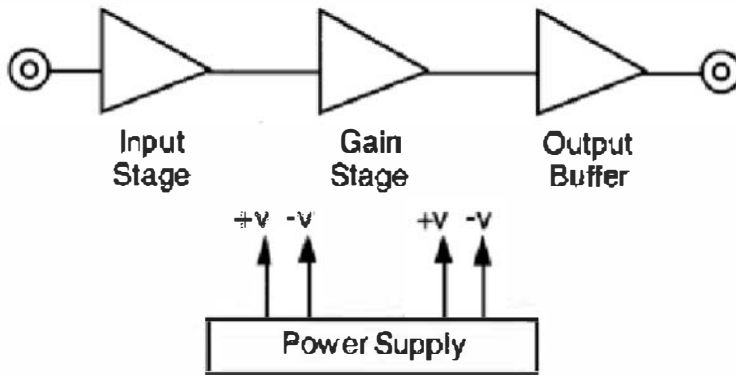


Fig4-9 Preamplifier block diagram

A preamplifier will also have a switch for selecting which source component is connected to your power amplifier for listening. This can be a series of relays controlled by electronic circuits, or simply a rotary mechanical switch.

The volume control is an important element of a preamplifier. Many preamps use a potentiometer, a continuously variable resistor in which a wiper attached to the preamp's volume knob moves in an arc around a resistive element. In preamps with remote control, this potentiometer (or "pot") is motor-driven. Another approach is to use an integrated-circuit chip that provides digital control over the analog signal. The preamp's front-panel volume control (and remote) sends a digital code to the chip that selects which resistors of a resistor-ladder network the audio is routed through. The resistor values determine the amount of attenuation. This chip-based attenuator provides stepped levels rather than a continuous volume adjustment. The steps are typically half a dB in the range where the control is most often used, and as much as 3dB at the very bottom of the range. An improvement on this technique is the stepped attenuator built from discrete resistors rather than the resistors within a chip. The discrete resistors can be switched into the circuit directly by the rotary volume knob, or by transistors outside the signal path. In the latter case, the volume control is merely an optical encoder that controls which FETs, and thus which resistors, are in the audio signal path. Fig.4-10 shows a cost-no-object implementation of the discrete-resistor stepped attenuator. Note that the attenuator has four large rotating elements rather than two. This is because the preamplifier in which it is used is fully balanced—a distinction explained later in this chapter. A rarely used technique is to employ a transformer with many secondary taps, with the volume control selecting the transformer tap.




Fig.4-10 A state-of-the-art switched-resistor attenuator. (Courtesy BAlabo and Pure Audio)

Most preamplifiers have separate miniature output stages to drive the tape output jacks. These output stages are nearly always of lower quality than the main output stage, but their sonic role is less crucial. Such products are said to have “independently buffered tape outputs.”

If you plan to drive two power amplifiers from your preamplifier (as in a bi-amplified system), make sure that the preamplifier’s output jacks are “independently buffered.” This means that each output is driven by its own output stage rather than a single output stage driving both outputs.

A feature missing from many high-end preamplifiers is a headphone jack. This connection lets you plug headphones directly into the preamp and listen to music without turning on your power amplifier. Most of the headphone amplifiers in preamps that do include this feature are of compromised quality. The best way to listen to headphones is through a separate headphone amplifier that accepts the main output from your preamplifier. (These products are

described in [Chapter 13](#).) A few preamplifiers, however, offer a high-quality headphone amplifier. If headphone listening is important to you, look for such a model.

The preamplifier's power supply delivers a supply of direct current (DC) to power the preamplifier's circuits. To prevent this DC from appearing at the preamplifier output along with the audio signal, DC blocking capacitors are often placed between stages, or at the preamplifier output. Some preamps avoid capacitors in the signal path by using a DC servo circuit that also prevents DC from appearing at the output jacks. A preamplifier without capacitors in the audio signal path is said to be direct-coupled.

Class A refers to a type of tube or transistor operation in which the tube or transistor is always biased on, meaning that it always conducts current. A discrete preamplifier is one that uses no integrated circuits (op-amps), only separate transistors or vacuum tubes. Discrete preamps are more expensive than op-amp-based designs. Virtually all state-of-the-art preamps are discrete Class A designs. (Descriptions of Class A operation are included in [Appendix B](#) and in [Chapter 5](#), "Power and Integrated Amplifiers.")

Balanced and Unbalanced Preamplifiers

As described earlier in this chapter, preamplifiers can have balanced inputs, balanced outputs, or both. Although all balanced preamps have XLR jacks, not all balanced preamplifiers are created equal. Two preamps that have balanced inputs and outputs can be very different in how they treat the signal.

Most preamps accepting a balanced signal immediately convert it to an unbalanced signal, perform the usual preamplifier functions (provide gain and volume adjustment) on the unbalanced signal, then convert the signal back to balanced just before the main output. A preamp with balanced inputs and outputs, but unbalanced internal topology, often adds two active stages to the signal path: a differential amplifier at the input and a phase splitter at the output. A differential amplifier converts a balanced signal to unbalanced; a phase splitter converts an unbalanced signal to balanced. This kind of preamplifier is shown in block form in [Fig.4-11](#).

The preferred, but much more expensive, method is to keep the signal balanced throughout the preamplifier. This technique requires double the audio circuitry; each portion of the balanced signal is amplified separately. Moreover, very close tolerances between halves of the balanced signal are required. The two signal halves must have identical gain and noise characteristics. Further, the four-element volume control must maintain precise tracking at all positions. Although the fully balanced circuit has more active devices in the signal path, the signal isn't subjected to a differential amplifier or a phase splitter. A fully balanced preamplifier is shown in block form in [Fig.4-12](#).

One way to tell if a preamplifier is fully balanced, or if it converts a balanced signal to unbalanced, then back to balanced, is to look at its volume control. If it has two elements (left and right channels), it has unbalanced internal circuitry. If the volume control has four elements (left +, left -, right +, right -), the preamp is fully balanced.

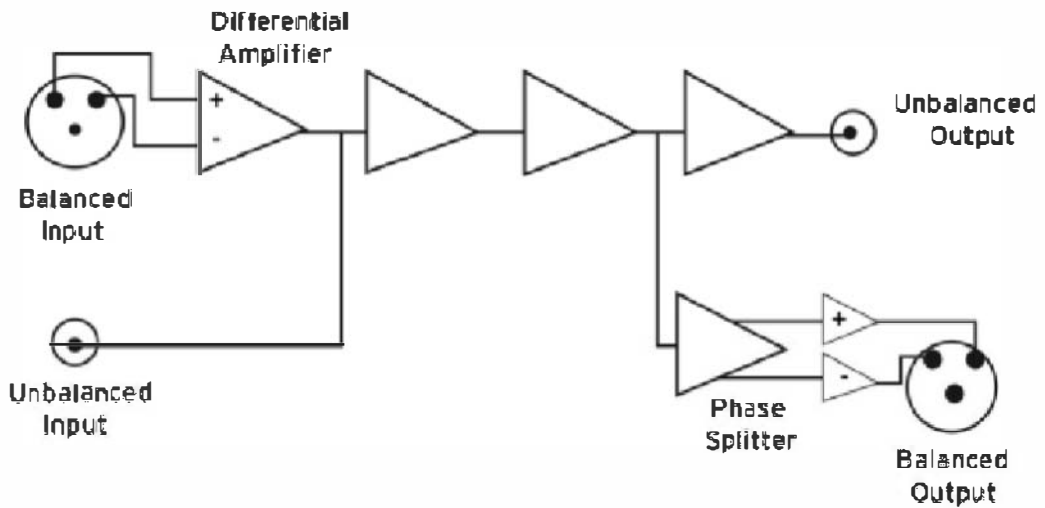


Fig.4-11 Some "balanced" preamplifiers convert a balanced input to unbalanced, then back to balanced.

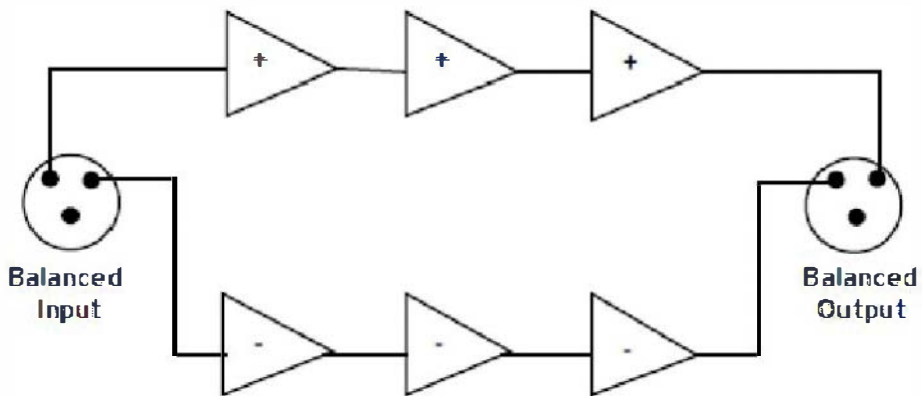


Fig.4-12 A fully balanced preamplifier keeps the signal balanced from input to output.

5 Power and Integrated Amplifiers

Introduction

The power amplifier—the last component in the signal chain before the loudspeakers—is the workhorse of a hi-fi system. It converts the low-level signal from the preamplifier into a signal powerful enough to drive the loudspeakers. It has a low-level input to receive the signal from the preamplifier, and terminals for connecting loudspeaker cables.

Power amplifiers range from \$400 solid-state stereo units to \$150,000/pair monoblocks. Choosing the right one for you begins with narrowing the range of possibilities. If you start by defining your requirements and setting your budget, the vast range of products will be quickly reduced to a more manageable number of contenders.

Because of the power amplifier's unique function, it differs from other components in size, weight, and use. High-quality power amplifiers are usually large and heavy. Moreover, the power amplifier is the component you don't need to touch and adjust. Consequently, power amplifiers are often placed on the floor near the loudspeakers, rather than in an equipment rack.

In this chapter we'll consider the wide range of power amplifiers, how to choose the amplifier best suited for your system, and explain the various amplifier technologies. Let's start by defining the terms associated with power amplifiers.

Stereo Power Amplifier: A power amplifier with two audio channels (left and right) on one chassis.

Monoblock Power Amplifier: A power amplifier with only one audio channel per chassis. Two monoblocks are required for stereo reproduction.

Multichannel Power Amplifier: An amplifier with, typically, five amplifier channels, for use in a home-theater system.

Watt: A unit of electrical power. Power is the ability to do work; in this case, the ability of the amplifier to make a loudspeaker's diaphragm move.

Power Output: The maximum amount of power the amplifier can deliver to the loudspeaker, measured in watts.

Load: The loudspeaker(s) the power amplifier must drive.

Tubed: A tubed power amplifier uses vacuum tubes to amplify the audio signal.

Solid-State: A solid-state power amplifier uses transistors to amplify the audio signal.

Single-Ended Amplifier: A power amplifier circuit topology in which the output device (a tube or transistor) always amplifies the entire musical waveform.

Class A: A mode of amplifier operation in which the amplifying device (tube or transistor) amplifies the entire musical waveform.

Class B: A mode of amplifier operation in which pairs of output devices (tubes or transistors) alternately "push" and "pull" current through the loudspeaker, with one output device amplifying the positive half of the waveform, another output device amplifying the negative half of the waveform. Also called a "push-pull" amplifier. (Compare with "Class A.")

Class A/B: A mode of amplifier operation in which the amplifier operates in Class A up to a small fraction of its output power, then switches to Class B.

Class D: A mode of amplifier operation in which the amplifier's output transistors turn fully on and fully off. Also called a "switching" amplifier.

Hybrid: A power amplifier combining vacuum tubes and solid-state devices. Usually, the input and driver stages are tubed and the output stage is solid-state.

Bridging: Converting a stereo power amplifier into a monoblock power amplifier. Some power amplifiers have a rear-panel bridging switch. Also called “strapping” into mono.

Bi-amplifying or bi-amping: Driving a loudspeaker’s midrange and treble units with one amplifier, the woofer with a second amplifier.

Tri-amplifying or tri-amping: Driving a loudspeaker’s woofer, midrange, and treble units each with its own amplifier.

Digital Power Amplifier: An amplifier that accepts only a digital rather than an analog audio signal. The digital signal is converted to a stream of pulses that turn on and off the output transistors.

How to Choose a Power Amplifier

Unlike most of the other components in your system, power amplifiers vary greatly in electrical performance. Consequently, choosing a power amplifier requires careful system matching for not only musical compatibility, but electrical compatibility as well. While any digital source will function in a system (even though it may not be musically ideal), some power amplifiers just won't work well with certain loudspeakers on a technical level. Choosing a power amplifier is thus a technical as well as an aesthetic decision that requires careful attention to system matching. We'll discuss these technical factors throughout this chapter.

Because a power amplifier's cost is often proportional to its output power, read this chapter's "How Much Power Do You Need?" section to select just the right power rating for your needs. Whatever you plan to spend on your system, the power amplifier should consume 20 to 25% of your total budget.

The first choice in power amplifiers—a stereo unit or a pair of monoblocks—will be decided by your budget. Monoblocks generally start at about \$2500 per pair. At this price level, a single stereo unit may make more sense; with only one case, power cord, and shipping carton, the manufacturer can put more of the manufacturing cost into better parts and performance. I advise against monoblocks if your amplifier budget is less than about \$4000. There may be exceptions to this figure, but it nonetheless offers a broad guideline. Many excellent stereo models, for example, cost upward of \$6000. A very popular price range for high-quality power amplifiers is \$800–\$2000, with the \$2000 models sometimes offering musical performance close to that of the most expensive amplifiers.

Monoblocks generally perform better than a single stereo model for several reasons. First, because the two amplifier channels are separate chassis, there is no chance of sonically degrading electrical interaction between channels. Consequently, monoblocks typically produce more convincing soundstages than do stereo units. Second, monoblocks have completely separate power supplies, even down to the power transformers: the left- and right-channel amplifier circuits don't have to share their electrical current source. This gives monoblocks the ability to provide more instantaneous current to the loudspeaker, all other factors being equal. Finally, most manufacturers put their cost-no-object efforts into monoblocks, which are often the flagships of their lines. If you want all-out performance and can afford them, monoblocks are the way to go.

Integrated Amplifiers

At the other end of the scale from monoblocks is the *integrated amplifier*, in which a preamplifier and a power amplifier are combined in the same case. Though the power output from integrated amplifiers is generally lower than that from separate power amplifiers, integrations are much less expensive, and ideal for budget to moderately priced systems. True high-end integrated amplifiers start at about \$350; top models can cost as much as \$25,000.

High-end integrated amplifiers have changed radically in the past few years. Once limited to low-powered units from European manufacturers, with idiosyncratic operation and nonstandard connectors, integrated amplifiers have finally come into their own. Leading high-end manufacturers have realized that an integrated amplifier makes sense for many music lovers. The cost and convenience advantages of integrations are compelling: they take up less space, are easier to connect, reduce the number of cables in your system, and can even offer the performance of separate components. Now that high-end manufacturers have taken the integrated amplifier seriously, they're putting their best technology and serious design efforts into such models.

Consequently, many manufacturers have enjoyed booming sales of integrated amps in the \$2000–\$4000 price range that produce 50–150Wpc (watts per channel) of power. Some manufacturers have even included a quality tuner with their integrated amplifier. Not so long ago, the term “high-end stereo receiver” was an oxymoron. Today, however, there's no reason a receiver designed and built with the care given to separate components should offer anything but high-end musical performance.

These newer integrated amplifiers have also overcome one of the limitations of earlier designs: the inability to upgrade only the power-amplifier or preamplifier section. Today's integrations often include preamplifier-out jacks for connecting the integration to a separate, more powerful amplifier. They also often have power-amplifier input jacks, for those who want to upgrade the preamplifier section. Higher-end integrated amplifiers feature a dual-mono design in which the left and right audio channels are completely separate from each other, even down to the power transformers. These premium-quality integrations also boast technologies found in upper-end preamplifiers, such as discrete-resistor stepped attenuators (described in the previous chapter), a discrete Class A input stage, and fully balanced operation.

The digital age has driven a radical transformation of the integrated amplifier from a simple, no-frills product into a technological showcase. For example, some of today's integrated amplifiers that include a built-in digital-to-analog converter also offer the same suite of *networking* and *wireless streaming* functions described in the previous chapter, on [page 58](#). The integrated amplifier can become a digital hub with Internet connectivity, multiple digital inputs, wireless streaming that includes the ability to stream music from a mobile device to the integrated amplifier, and can also function as an integral part of a networked audio system.

Other useful features of integrated amplifiers include *gain offset* and *theater bypass*. Gain offset allows you to individually attenuate the signal level for each input, to compensate for the varying output levels of source components. This adjustment prevents large jumps in playback volume when switching between sources. Theater bypass is important if you plan to use the integrated amplifier as part of a home-theater system. This switch, or sometimes a dedicated input, sets the integrated amplifier at a fixed gain (the volume control is disabled) so that you can calibrate the channel levels with an A/V controller and maintain that calibration even as you change the overall volume level of all channels in a multichannel system. [Chapter 12](#) includes a complete description of how a stereo preamplifier or stereo integrated amplifier works with an A/V controller in a home-theater system.

When choosing an integrated amplifier, combine the advice in [Chapter 4](#) ("Preamplifiers") with the guidelines in the rest of this chapter. If your budget is under \$5000 for amplification, seriously consider one of the new breed of high-end integrated amplifiers rather than separates.

How Much Power Do You Need?

This is the first question to answer when shopping for a power amplifier or integrated amplifier. Power output, measured in watts into a specified loudspeaker impedance, varies from about 20Wpc in a very small integrated amplifier to about 1000Wpc in the largest monoblocks. Most high-end power amplifiers put out between 80 and 250Wpc. Single-ended-triode amplifiers, described later in this chapter, generally produce between 3 and 20Wpc.

Choosing a range of amplifier power output that's appropriate for your loudspeakers, listening tastes, room, and budget is essential to getting the best sound for your money. If the amplifier is underpowered for your needs, you'll never hear the system at its full potential. The sound will be constricted and fatiguing, will lack dynamics, and the music will have a sense of strain on climaxes. Conversely, if you spend too much of your budget on an amplifier bigger than you need, you may be shortchanging other components. Choosing just the right amplifier power is of paramount importance.

The amount of power needed varies greatly according to loudspeaker *sensitivity*, loudspeaker *impedance*, room size, room acoustics, and how loudly you like to play music. Loudspeaker sensitivity specifies how high a sound-pressure level (SPL) the loudspeaker will produce when driven by a certain power input, and is, by far, the biggest determining factor in choosing an appropriate power output. A typical sensitivity specification will read "88dB SPL, 1W/1m" (often shortened to "88dB/W/m"). This means that the loudspeaker will produce an SPL of 88 decibels (dB) with one watt of input power when measured at a distance of one meter. (The decibel is explained in detail in [Appendix A](#), "Sound and Hearing.") Although 88dB is a moderate listening volume, a closer look at how power relates to listening level reveals that we need much more than 1W for music playback.

Each 3dB increase in sound-pressure level requires a doubling of amplifier output power. Thus, a loudspeaker with a sensitivity of 88dB at 1W would produce 91dB with 2W, 94dB with 4W, 97dB with 8W, and so on. For this loudspeaker to produce musical peaks of 109dB, we would need an amplifier with 128W of output power.

Now, say we had a loudspeaker rated at 91dB at 1W/1m—only 3dB more sensitive than the first loudspeaker. We can quickly see that we would need only half the amplifier power (64W) to produce the same volume of 109dB SPL. A

loudspeaker with a sensitivity of 94dB would need just 32W to produce the same volume. A speaker with higher sensitivity can simply convert more of the amplifier's power into sound.

This relationship between amplifier power output and loudspeaker sensitivity was inadvertently illustrated in an unusual demonstration more than 65 years ago. In 1948, loudspeaker pioneer Paul Klipsch conducted a demonstration of live vs. reproduced sound with a symphony orchestra and his Klipschorn loudspeakers. His amplifier power: 5W to 30W (accounts differ). The Klipschorns are so sensitive (an astounding 104dB/W/m) that they will produce very high volumes with very little amplifier power. Klipsch was attempting to show that his loudspeakers could closely mimic the tonal quality and loudness of a full symphony orchestra.

The other end of the speaker-sensitivity spectrum was illustrated by a demonstration I attended of an exotic new loudspeaker. During the demo, the music was so quiet that I could barely hear it. I looked at the power amplifiers—300Wpc monsters with large power meters—and was astonished to see that the power meters were nearly constantly pegged at full power. This unusual speaker converted only a minuscule amount of the amplifier's output power into sound.

The importance of loudspeaker sensitivity is also demonstrated by today's 3Wpc single-ended-triode amplifiers, which can produce moderately loud listening levels through high-sensitivity speakers. These examples of huge variations in sound-pressure level and amplifier power illustrate how loudspeaker sensitivity greatly affects how big an amplifier you need. Even a small difference in loudspeaker sensitivity—2dB, say—changes your amplifier power requirements.

The dBW Power Rating

We've seen that every doubling of amplifier power yields a volume increase of 3dB. Consequently, there is a 3dB difference between a 10W amplifier and a 20W amplifier, but also between a 500W amplifier and a 1000W amplifier. Although the output power is vastly greater between 500 and 1000W than between 10 and 20W, the difference in the volume level produced is still only 3dB. That's why we must consider the ratio of output powers, not the number of watts difference, when comparing amplifier power ratings.

Fortunately, there is a way of expressing amplifier power that takes into account this exponential relationship between dB and power, a rating called the *dBW*. The "W" after "dB" means that 0dB is referenced to 1W of output power. A 10W amplifier has a dBW rating of 10dBW, a 20W amplifier has a dBW rating of 13dBW—twice the power, or the ability to play 3dB louder than a 10W amplifier. A 40W amplifier is rated at 16dBW. One hundred watts equals 20dBW, and a 1000W amplifier is said to have 30dBW of output power. The dBW rating expresses the ratio of amplifier output power referenced to 1W of output power. The relationship between power output (into an 8 ohm load) and dBW is shown in the following table.

Watts	dBW	Watts	dBW	Watts	dBW
1.0	0	10.0	10	100	20
1.25	1	12.5	11	125	21
1.6	2	16	12	160	22
2.0	3	20	13	200	23
2.5	4	25	14	250	24
3.2	5	32	15	320	25
4.0	6	40	16	400	26
5.0	7	50	17	500	27
6.3	8	63	18	630	28
8.0	9	80	19	800	29

Most manufacturers don't include dBW ratings in their amplifiers' specifications. You can, however, perform this conversion yourself with a pocket calculator that has a "log" button. Punch in the number of watts you want to convert to dBW, find the logarithm (base 10) of that number by pressing the log key, then multiply by 10. This will give you the dBW value when driving an 8

ohm load. To find the dBW rating with a 4 ohm load, subtract 3dB from the result. To calculate for 2 ohms, subtract 6dB from the result.

For example, if an amplifier puts out 150W into 8 ohms, enter “150” into the calculator and press the log key; “2.176” will appear in the display. Then press “x” and “10” to multiply the log by 10, to get the dBW rating of 21.76dBW.

If this same amplifier puts out 225W into 4 ohms, enter “225” into the calculator, press log to get 2.352, then multiply by 10 to get the intermediate result of 23.52dBW. But this is a 4 ohm power rating; subtracting 3 from the result will give you the final dBW rating of 20.52W into 4 ohms.

If you’re choosing between 60 and 110Wpc amplifiers, you can calculate the difference in dB between them. Converting 60W (at 8 ohms) into dBW gives us 17.78dBW. Converting 110W (at 8 ohms) into dBW works out to 20.41dBW. The difference is a little less than 3dB. Note that the same difference in dB exists between 600 and 1100W amplifiers as between 6 and 11 W amplifiers.

The concept of dBW becomes even more useful when we take a closer look at amplifier output ratings. Most of the power figures we’ve discussed are with the amplifier driving an 8 ohm load. But most loudspeakers have impedances of less than 8 ohms, and some dip below 2 ohms. How does our power amplifier behave when asked to drive low-impedance loudspeakers? (Impedance is explained in [Chapter 6](#), “Loudspeakers,” and in [Appendix B](#), “Audio and Electronics Basics.”)

Instead of considering only the amplifier’s 8 ohm power rating, look at how the power output changes into 4 and 2 ohm loads. An amplifier acting as a perfect voltage source will double its power every time the load impedance is cut in half; this amplifier would be capable of driving 100W into 8 ohms, and could deliver 200W into 4 ohms, 400W into 2 ohms, or 800W into 1 ohm. (A more technical discussion of this aspect of amplifier performance appears in [Chapter 16](#).)

Some amplifiers barely increase their output power when driving 4 ohms; others can double it. This means that not all “100Wpc” amplifiers are created equal. One “100Wpc” amplifier might put out 150W into 2 ohms, while another might deliver 400Wpc into 2 ohms. This ability to drive low-impedance loads (specifically, to deliver lots of current) has a large influence on an amplifier’s sound and subjective power capability. Loudspeakers have dips in their impedances at certain frequencies, which puts greater current-delivery demands on the power amplifier.

This is where our dBW rating again comes into play. When calculating dBW, subtract 3dBW from the 8 ohm dBW rating when the amplifier is driving 4 ohms, and subtract 6dBW when the amplifier is driving 2 ohms. In the example of the

two “100Wpc” amplifiers in the preceding paragraph, each is rated at 100Wpc into 8 ohms, or 20dBW. But the first amplifier puts out only 150W into 2 ohms, while the second puts out 400W into 2 ohms. Converting the 2 ohm power ratings into dBW, we see that the first amplifier has a dBW rating of 15.76dBW at 2 ohms, while the second has a dBW rating of 20dBW into 2 ohms. Quite a difference—more than 4dB—between two amplifiers that are both rated at 100Wpc into 8 ohms. This is why not all “100W” amplifiers are created equal.

Why Amplifier Output Current Matters

This difference in output power into low-impedance loudspeakers has real-world consequences. The ability to increase output power into low impedances indicates how much current the amplifier can deliver to the loudspeaker. It is current flow through the loudspeakers' voice coils (in dynamic loudspeakers) that creates the electromagnetic force that causes the cones and domes to move, and thus produce sound. If current flow through the voice coil is constrained, so is the music.

In the previous chapter, I described how source components such as DACs have an easier time driving a preamplifier of high input impedance than one of low input impedance. I likened the situation to a water faucet and hose: the water pressure is voltage, the flow of water through the hose is electrical current, and squeezing the hose forms a resistance (impedance) to the flow. You can extend this analogy to a power amplifier driving a loudspeaker, with the loudspeaker's impedance forming the resistance in the hose. The lower the loudspeaker's impedance, the less the resistance to current flow from the amplifier, and the harder the amplifier must work to deliver current to the loudspeaker. As we'll see in [Appendix A](#), if the impedance is halved (say, from 8 to 4 ohms), the amplifier is asked to deliver twice the amount of current to the loudspeaker.

If the amplifier isn't up to the job, the musical result is strain or even distortion on musical peaks, weak bass, loss of dynamics, hardening of timbre, and a collapsing sound-stage. In short, we can hear the amplifier give up as it runs out of power. Conversely, amplifiers that can continue increasing their output power as the impedance drops generally have very deep, extended, and powerful bass, virtually unlimited dynamic headroom, a sense of ease and grace during musical peaks, and the ability to maintain accurate timbre and soundstaging, even in loud passages. If you have relatively high-impedance loudspeakers with no severe impedance dips, you're much less likely to encounter sonic problems, even with modest power amplifiers—the loudspeaker simply demands less current.

Amplifiers with high current capability (indicated by their ability to increase output power into low impedances) are often large and expensive. Their current capability comes from massive power transformers, huge power supplies, and lots of output transistors—all expensive items. [Fig.5-1](#) shows such a power amplifier. The transformer and filter capacitors visible in the middle of the amplifier can supply large amounts of current to the circuits, and the output stage includes many transistors working together for increased current capacity. To keep them

running relatively cool, these transistors are mounted on large heatsinks that transfer heat generated by the transistors to fins on the amplifier's exterior, so that it can be dissipated into the air.

In the "Specifications and Measurements" section on power amplifiers in [Chapter 16](#), we'll look at how to evaluate a power amplifier's ability to deliver current into low-impedance loads. You can also figure out for yourself which amplifiers are good at driving low-impedance speakers by converting their output-power ratings at 8 and 4 ohms into dBW. The less the difference in the dBW ratings into 8 and 4 ohms, the better the amplifier's ability to deliver current.

Keep in mind, however, that not all audio systems require large power amplifiers. If your loudspeakers are sensitive and have a fairly high impedance, their current demands are vastly lower. Consequently, smaller amplifiers should work just fine. A system that includes a powered subwoofer will place lower power demands on the power amplifier or integrated amplifier. Because the amplifier needn't drive a loudspeaker's woofer at the lowest frequencies, it can be of lower power. And as we'll see in the discussion of single-ended-triode amplifiers later in this chapter, amplifiers that produce as little as 3Wpc and have a very limited ability to deliver current can still sound highly musical.




Fig.5-1 Inside a high-current power amplifier (Courtesy BALABO and Pure Audio)

What to Look For when Comparing Power Ratings

When comparing amplifier power ratings, make sure the specified power is continuous or *RMS* rather than peak. Some manufacturers will claim a power output of 200W, for example, but not specify whether that power output is available only during transient musical events such as drum beats, or if the amplifier can deliver that power continuously into a given load. RMS stands for “root mean square,” a mathematical calculation expressing the effective, or average, power output. These days, however, very few amplifiers are specified by peak power.

Another way manufacturers exaggerate power ratings is by not specifying the *power bandwidth*. This term describes the frequency range over which a power amplifier can deliver the power output specified. A power amplifier delivering 200W only at 1kHz is far less powerful than one that can deliver 200W over the full audio bandwidth of 20Hz–20kHz, the general range of human hearing. You’ll often see mass-market audio/video receivers with power-output ratings specified only at 1kHz, or from 50Hz to 20kHz. Further, stereo power amplifiers can deliver more power when only one channel is driven—look for the words “both channels driven.” The maximum power output should also be specified at a certain maximum level of distortion.

You can see the potential for misleading claims of power-amplifier output. At one time, the abuses were so bad that the Federal Trade Commission (FTC) stepped in to regulate power claims—the only example of an audio specification being regulated by a governmental body. The FTC mandate for power ratings requires that the power rated be continuous (not peak), that the load impedance and bandwidth be specified, and that the total harmonic distortion (THD) be given at full power and measured over the audioband. You may see a power specification that reads “50Wpc continuous (or RMS) power into 8 ohms, both channels driven, 20Hz–20kHz, with less than 0.1% THD.” A power specification including all these conditions is called an “FTC power rating.” Some manufacturers no longer adhere to the FTC-mandated power ratings, figuring that the issue has blown over and is no longer enforced. You see fudged power ratings for mass-market A/V receivers that must now power five or seven loudspeakers rather than two, and in single-ended-triode amplifiers that can’t meet the FTC’s stringent requirements for power-output specifications.

If you’re amplifier-shopping for low-impedance loudspeakers, look at the power-output specifications into 4 ohms. Make sure you see the words

continuous or *average* in the power rating, and that the bandwidth and distortion are specified. These figures don't tell us what we need to know about the amplifier's musical qualities, but they at least indicate good technical performance.

Why Amplifier Power Isn't Everything

We've seen how loudspeaker sensitivity greatly affects how much amplifier power you need, and how power amplifiers with the same power rating into 8 ohms can differ radically in their abilities to drive loudspeakers. Now let's look at some other factors influencing how much amplifier power you need.

The first is room size. The bigger the room, the more amplifier power you'll need. A rough guide suggests that quadrupling the room volume requires a doubling of amplifier power to achieve the same sound-pressure level. How acoustically absorptive or reflective your listening room is will also affect the best size of amplifier for your system. If we put the same-sensitivity loudspeakers in two rooms of the same size, one room acoustically dead (absorptive) and the other acoustically live (reflective), we would need roughly twice the amplifier power to achieve the same sound-pressure level in the dead room as in the live room.


Finally, how loudly you listen to music greatly affects how much amplifier power you need. Chamber music played softly requires much less amplifier power than rock or orchestral music played loudly. The relationships between loudspeaker sensitivity, room size, room acoustics, and amplifier power are shown in [Fig.5-2](#).

We can see that a low-sensitivity loudspeaker, driven by orchestral music in the large, acoustically dead room of someone who likes high playback levels, may require many times the amplifier power needed by someone listening to chamber music at moderate listening levels through high-sensitivity loudspeakers in a small, live room. A 20Wpc amplifier may satisfy the second listener; the first listener may need 750Wpc.

POWER NEEDED IN WATTS (dBW)

160 (22)
140 (21.5)
120 (20.75)
100 (20)
80 (19)
60 (17.75)
40 (16)
20 (13)

LIVE ROOM




LOUDSPEAKER SENSITIVITY
(FOR 1-WATT INPUT MEASURED AT 1 METER)

POWER NEEDED IN WATTS (dBW)

160 (22)
140 (21.5)
120 (20.75)
100 (20)
80 (19)
60 (17.75)
40 (16)
20 (13)


AVERAGE ROOM



LOUDSPEAKER SENSITIVITY
(FOR 1-WATT INPUT MEASURED AT 1 METER)

160 (22)
140 (21.5)

DEAD ROOM



LOUDSPEAKER SENSITIVITY
(FOR 1-WATT INPUT MEASURED AT 1 METER)

Fig.5-2 The relationship between loudspeaker sensitivity, room size, room acoustics, and power-amplifier output power (after *High Fidelity*)

Other Power-Amplifier Considerations

Tubes vs. Transistors

In the late 1950s, when transistor amplifiers were introduced, it appeared that the days of the vacuum tube were over. Transistors were smaller, lighter, and cheaper than tubes, ran cooler, and produced more output power. As if that weren't enough, a transistor amplifier didn't need an output transformer, a component that added considerably to an amplifier's size, weight, and cost. Virtually overnight, nearly all audio manufacturers scrapped their tubed amplifiers in favor of transistor models.

But many music lovers found the sound of these newfangled "solid-state" amplifiers unlistenable. They likened the sound to that of a pocket transistor radio, only louder. Unfortunately, following the introduction of transistor amps, those perceptive audiophiles couldn't buy a new tubed amplifier.

The modern tubed amplifier was created in 1970 by William Zane Johnson, of Audio Research Corporation. Like many music lovers, he found the sound of tubed amplifiers more musical. Johnson's demonstration of a tubed amplifier at a 1970 hi-fi show prompted one showgoer to remark, "You've just set the audio industry back ten years!" But instead of being a setback, Johnson's amplifier began a renaissance in tubed equipment that is still going strong more than 45 years later.

The perennial tubes-vs.-transistors debate arises when you're faced with choosing a power amplifier. Tubed models can offer stunning musical performance, but they have their drawbacks. Here are the advantages and disadvantages of tubed power amplifiers. (For the moment, I'll confine my observations to conventional "push-pull" tubed amplifiers, not the exotic single-ended-triode varieties described later in this chapter.)

Tubed power amplifiers are more expensive than their similarly powered solid-state counterparts. The cost of tubes, transformers, and more extensive power supplies all make owning a tubed power amplifier a more expensive proposition than owning a solid-state design. Moreover, the tubes will need replacing every few years, further adding to the cost of ownership. A set of sixteen EL34 output tubes, for example, can cost from \$300 to \$1000.

In terms of bass performance, tubed power amplifiers can't compete with good solid-state designs, all else being equal. Tubes provide less control of the

bass, resulting in a sound that's less punchy, taut, and extended. Further, tubed power amplifiers often have limited current delivery into low-impedance loads, making them a poor choice for current-hungry loudspeakers. Tubes also require monthly biasing (small adjustments made with a screwdriver) to maintain top performance. Biasing is very easy, but some music lovers would prefer not to have to think about performing routine maintenance on their music-playback systems. A few tubed amplifiers, however, offer automatic-biasing circuits.

Power-amp tubes can suddenly fail, sometimes in smoke and (momentary) flames. Such dramatic failures are rare, however; I've used tubed amplifiers for decades and in that time have had two tubes fail, both uneventfully.

Finally, there is the possibility of small children or pets burning themselves on hot tubes. Many, but not all, tubed power amplifiers have exposed output tubes. If you have small children, consider an amplifier whose tubes are surrounded by a ventilated metal cage.

Given these drawbacks, why would anyone want to own a tubed power amplifier? It's simple: tubes can sound magical. When matched with an appropriate loudspeaker, tubed power amplifiers offer exceptional musicality, in my experience. Even small, moderately priced tubed amps have more than a taste of tube magic.

Many important aspects of music reproduction seem to come naturally to tubed amplifiers. They generally have superb presentation of instrumental timbres, smooth and unfatiguing treble, and spectacular soundstaging. The hard, brittle, edgy midrange and treble of many solid-state amplifiers is contrasted with the purity of timbre and sense of ease conveyed by a good tubed amp. Music has a warmth, ease, and natural musicality when reproduced by many tubed designs. The soundstage has an expansive quality, with a sense of bloom around instrumental images. This isn't to say there aren't good-sounding solid-state power amplifiers, only that tubes seem to more consistently deliver the musical goods. A tubed amplifier's softer bass is often willingly tolerated for its magical midrange, treble, and sometimes holographic soundstaging.

Beware, however, of the tubed power amplifier that makes every recording sound lush; such an amp is probably imposing its own sonic signature on the music. Some tubed amplifiers have euphonic colorations, or departures from accuracy that are initially pleasing to the ear. Using several different recordings, listen for a change in timbre, spatial presentation, or bass quality. These characteristics should change with the recording; if they don't, it's likely that the power amplifier is imposing a common sonic signature on all of them. Ideally, the

power amplifier will reveal exactly what's on the recording while adding no characteristics of its own.

As good as tubed amplifiers can sound, solid-state amplifiers have some decided sonic and technical advantages. For example, tubed models are no match for solid-state amplifiers in bass performance. Transistor power amplifiers have bass that is tighter, deeper, and much more solid than tubed amps. The feeling of bass tautness, kick, extension, and power are all better conveyed by solid-state amplifiers, regardless of how good the tubed amplifier is. Speaking technically, solid-state amplifiers can deliver more current to low-impedance loudspeakers, making them a better choice for such speakers.

No one but you can decide if a tubed power amplifier is ideal for your system. I strongly suggest, however, that you audition at least one tubed amp before making a purchasing decision. You may get hooked.

Balanced Inputs

If your preamp has balanced outputs, you may want to consider a power amplifier with balanced inputs. Most power amps with balanced inputs also provide unbalanced inputs, allowing you to compare these two connection methods before deciding which to leave in your system. Some combinations of preamp and power amp sound better via balanced connection, others via the unbalanced jacks. The best way to discover which method is better is by listening to both. (See [Chapter 4](#) for a discussion of why some preamplifiers may sound better from their unbalanced outputs.) Borrow a pair of interconnects from your dealer so you can make this comparison. If you have a power amplifier with both balanced and unbalanced inputs, but use the unbalanced inputs, be sure to insert a shorting plug in the amplifier's balanced input jacks. This is a small, U-shaped plug that shorts pin 3 of the balanced input connector to ground, preventing noise from getting into the amplifier.

Bridging

Some stereo power amplifiers can be “bridged” to function as monoblocks. *Bridging* configures a stereo amplifier to function as a more powerful single-channel amplifier. The amplifier will have a switch on the rear panel to convert it to bridged operation. Note that two bridged amplifiers are needed for stereo. If you have a stereo amplifier that can be bridged and you want more power, simply buy a second, identical amplifier and bridge the two for more total power. In theory, bridging results in a fourfold increase in output power. That’s because bridging doubles the amplifier’s maximum output voltage and, according to Ohm’s law, quadruples the power. In practice, however, bridging roughly doubles an amplifier’s power rating into a 4 ohm load, due to the amplifier’s current-output limitations.

Bridging changes the amplifier’s internal connections so that one channel amplifies the positive half of the waveform and the other channel amplifies the negative half. The loudspeaker is connected as the “bridge” between the two amplifier channels, instead of between one channel’s output and ground.

Bridging is most beneficial when the power amplifiers are asked to drive low-sensitivity, high-impedance (8 ohms nominal) loudspeakers. High-impedance speakers are driven more by voltage than by current. Conversely, low-impedance speakers demand more current from the power amplifier. Bridging doubles the amplifier’s maximum output voltage, but quadruples its maximum current output (because two amplifier channels are now driving one loudspeaker). Moreover, connecting a 4 ohm speaker to a bridged power amplifier causes the amplifier to “see” a 2 ohm load, further stressing the amplifier’s current capacity. The result can be amplifier overheating, which will either damage the amplifier, or activate its protection circuit and shut down the amplifier while music is playing.

Bi-Amping

Any loudspeaker with two sets of input terminals is a candidate for *bi-amping*. In a bi-amped setup, two power amplifiers are used to drive one loudspeaker. One amplifier drives the woofer, the other drives the midrange and tweeter.

Let's first consider passive bi-amping. The preamplifier output is split into two left-channel signals and two right-channel signals. (Many preamps have two stereo output pairs for this purpose.) The left-channel preamp output drives the left-channel inputs on both power amps. The right-channel preamp output drives the right-channel inputs on both power amps. One power amp drives the bass input on both loudspeakers, the other power amp drives the midrange and tweeter input on both loudspeakers. Note that the loudspeaker's crossover still operates; both amplifiers receive and amplify a full-bandwidth signal. This is called passive bi-amping because the division into bass and midrange/treble frequencies is performed by passive components (capacitors, resistors, and inductors) in the loudspeaker's crossover.

In active bi-amping, rather than allow the loudspeaker's crossover to split the audioband, an electronic crossover between the preamp and power amps does the job. Active bi-amping divides the audioband before the power amplifiers, not after, as in passive bi-amping. The losses and problems inherent in any passive crossover are eliminated by active bi-amping. Moreover, each power amplifier receives only the portion of the audioband it will amplify. Active bi-amping is more efficient than passive bi-amping, but requires bypassing the loudspeaker crossovers and adding an electronic crossover.

Bi-amping not only increases the total power delivered to the loudspeaker, it allows each amplifier to work over a narrower frequency range. We can thus choose a bass amplifier for its superior bottom-end performance, regardless of how it sounds in the mids and treble, and a midrange/treble amplifier without concern for its bass. An ideal combination may include a large solid-state amplifier for the bass and a sweet-sounding tubed amplifier for the top end. The small tubed amplifier is relieved of the burden of driving the bass, allowing it to do what it does best. Bi-amping is one way of getting the best of both the tubed and solid-state worlds.

Another advantage of active bi-amping is that the woofer amplifier can be connected directly to the woofer's voice coil rather than through a passive crossover network, particularly the series inductor found in most woofer

crossovers. This provides much tighter damping (the ability of the amplifier to control woofer motion) and results in tighter, deeper bass reproduction.

When bi-amping is chosen, the two amplifiers' gains (their amounts of amplification) must be exactly matched. If you're using identical amplifiers, the gain will be inherently matched. But with different amplifiers, a gain mismatch will produce too much or too little bass. An amplifier can be padded down at the input with an attenuator to reduce the input level and match its output to the second amplifier. If you're using an active external crossover, it may include independent output-level controls to balance the levels between the woofer and midrange/tweeter amplifiers.

Bi-amping can be tricky. I don't recommend it for most systems unless you fully understand what's happening, or unless your dealer helps you. One serious drawback is that replacing the speaker's internal crossover with an external one changes the crossover frequencies and characteristics the loudspeaker designer decided were optimal for that speaker's drivers. Many loudspeaker designers don't take kindly to someone second-guessing their crossover characteristics, and active bi-amping may void the speaker's warranty.

A few manufacturers—Linn Products, for example—have designed into their products easy bi-amping upgradeability; bi-amping is just a matter of installing active crossover boards in the company's power amplifiers, adding a second amplifier, and removing the crossovers inside the loudspeaker. In this case, the loudspeaker works optimally with the external crossover.

Incidentally, adding an active subwoofer (one with an integral power amplifier; see [Chapter 6](#)) automatically makes your system bi-amped. The subwoofer's internal amplifier drives the low bass, and your main amplifier handles only the upper bass, mids, and treble. Installing an active subwoofer is one way of getting the benefits of bi-amping without changing amplifiers or adding another electronic crossover. Loudspeakers that come with a dedicated woofer amplifier are also inherently bi-amplified.

The benefits of bi-amping can be extended to tri-amping, in which three amplifiers drive one loudspeaker.

What to Listen For

How can you tell if the power amplifier you're considering will work well with your loudspeakers? Simple: Borrow the amplifier from your dealer for the weekend and listen to it. This is the best way not only to assess its musical qualities, but to determine how well it drives your loudspeakers. In addition, listening to the power amplifier at home will let you hear if the product's sonic signature complements the sound of the rest of your system.

The next best choice is if the dealer sells the same loudspeakers you own, and allows you to audition the combination in the store. If neither of these options is practical, consider bringing your loudspeakers into the store for a final audition.

All the sonic and musical characteristics described in [Chapter 3](#) apply to power amplifiers. However, some sonic characteristics are more influenced by the power amplifier than by other components.

The first thing to listen for is whether the amplifier is driving the loudspeakers adequately. Weak bass is a sure indication that the amplifier is underpowered for a particular pair of speakers. If the bass is soggy, slow, or lacks punch, the amplifier probably isn't up to the job of driving your loudspeakers. Other telltale signs that the amplifier is running out of current include loss of dynamic impact, a sense of strain on musical peaks, hardening of timbre, reduced pace and rhythm, and soundstage collapse or congestion. Let's look at each of these individually.

First, play the system at a moderate volume. Select music with a wide dynamic range—either full orchestral music with a loud climax accompanied by bass drum, or music with a powerful rhythmic drive from bass guitar and kick drum working together. Audiophile recordings typically have much wider dynamic ranges than general-release discs, making them a better source for evaluation. Music that has been highly compressed to play over the radio through a 3" car speaker will tell you less about what the system is doing dynamically.

After you've become used to the sound at a moderate level, increase the volume—you want to push the amplifier to find its limits. Does the bass seem to give out when you turn it up, or does the amplifier keep on delivering? Listen to the dynamic impact of kick drum on a recording with lots of bottom-end punch. It should maintain its tightness, punch, quickness, and depth at high volume. If it starts to sound soggy, slow, or loses its power, you've gone past the point at which

the amplifier can operate comfortably. After a while, you can get a feel for when the amplifier gets into trouble. Is the sound strained on peaks, or effortless?

NOTE: When performing this experiment, be sure not to overdrive your loudspeakers. Turn down the volume at the first sign of loudspeaker overload (distortion or a popping sound).

Compare the amplifier's sound at high and low volumes. Listen for brass instruments becoming hard and edgy at high volumes. See if the soundstage degenerates into a confused mess during climaxes. Does the bass drum lose its power and impact? An excellent power amplifier operating near its maximum power capability will preserve the senses of space, depth, and focus, while maintaining liquid instrumental timbres. Moreover, adequate power will produce a sense of ease; lack of power often creates listener fatigue. Music is much more enjoyable through a power amplifier with plenty of reserve power.

All the problems I've just described are largely the result of the amplifier running out of current. Just where this happens is a function of the amplifier's output power, its ability to deliver current into the loudspeaker, the loudspeaker's sensitivity and impedance, the room size, and how loudly you expect your hi-fi system to play. Even when not pushed to its maximum output, a more powerful amplifier will often have a greater sense of ease, grace, and dynamics than a less powerful amplifier.

Assuming that the power amplifier can drive your loudspeakers, let's evaluate its musical characteristics.

First, listen for its overall perspective. Determine if its sound is forward or laid-back, then decide if the amplifier is too "up-front" or too recessed for your system or tastes. A forward perspective seems to put you in Row C of the concert hall; a laid-back perspective puts you in Row W. Forwardness gives the impression of presence and immediacy, but can quickly become fatiguing. A sound that's too laid-back is less exciting and compelling. If the amplifier's overall perspective isn't right, whatever else it does well may not matter.

Next, listen for brightness, hardness, treble grain, and treble etch. If you feel a sense of relief when the volume is turned down, you can be pretty sure that the treble is at fault. Listen to midrange timbres, particularly of voices and the violin, piano, and flute. Midrange textures should be sweet and liquid, and devoid of grain, hardness, or a glassy edge. Many power amplifiers—particularly inexpensive solid-state models—make the treble dry, forward, and unpleasant. If you hear even a hint of these characteristics in a few hours' listening, rest assured

that they will grow only more annoying over time. Many moderately priced amplifiers don't add these unmusical characteristics—keep shopping.

Power amplifiers vary greatly in their ability to throw a convincing soundstage. With naturally miked recordings, the soundstage should be deep, transparent, and focused, the loudspeakers and listening room seeming to disappear. Instruments should hang in space within the recorded acoustic. The sound should provide a transparent, picture-window view into the music. If the sound is flat, veiled, or homogenized, you haven't found the right amplifier.

Bass performance is also largely influenced by the power amplifier. It should sound taut, deep, quick, dynamic, and effortless. The music should be propelled forward, involving your body in its rhythmic drive. Kick drum should have depth, power, and a feeling of suddenness. The bass guitar and kick drum should provide a tight, solid, and powerful musical foundation.

The best way to find out which power amplifiers sound good and are well made is to confer with a trusted dealer and read product reviews in reputable high-end magazines. Responsibly written reviews can save you from spending time and trouble on unworthy contenders, and point you toward products that offer promise. Thorough reviews will also include insightful commentary on an amplifier's technical performance. If the amplifier has technical characteristics that make its use with a particular type of loudspeaker questionable or unsuitable, the review will reveal this fact. A review's technical section may state that, "Although the XYZ power amplifier sounded terrific with the Dominator 1 loudspeakers, its high output impedance and limited current delivery make it a poor choice with low-sensitivity loudspeakers (below 86dB/W/m), or those with an impedance below 5 ohms." This kind of commentary can point you toward the best amplifiers for your loudspeakers.

In short, read lots of product reviews before drawing up a list of candidates to audition. Restrict candidates to those that fall within your budget, and take the few contenders home for an evening's listen. And don't be in a hurry to buy the first amplifier you audition—listen to several before choosing. Your effort and patience will be rewarded with a more musical amplifier, and one better suited to your system.

A Survey of Amplifier Types

So far I've described the types of mainstream power amplifiers that most audiophiles are likely to buy. But a number of variations on these basic designs are significant forces in the amplifier marketplace. These include the single-ended-triode amplifier, the single-ended solid-state amplifier, switching power amplifiers, and digital amplifiers.

Single-Ended-Triode Amplifiers

One of the most interesting trends in high-end audio over the past 25 years has been the return of the *single-ended-triode* power amplifier. The single-ended-triode (SET) was the first audio amplifier ever developed, dating back to Lee De Forest's patents of the Audion (1907) and triode amplifier (1912). SET amplifiers generally deliver very low power, sometimes just a few watts per channel.

You heard right: Large numbers of audiophiles are flocking to replace their modern power amplifiers with amplifiers based on a 100-year-old technology. Have the past 100 years of amplifier development been a complete waste of time? A surprising number of music lovers and audio designers think so.

The movement back to SET amplifiers began in Japan in the 1970s, specifically with designer Nobu Shishido, who combined SET amplifiers with high-sensitivity, horn-loaded loudspeakers. Many who heard SET amplifiers were startled by their goose-bump-raising musical immediacy and ability to make the music "jump" out of the loudspeakers. Thus began the rage for SET amplifiers in Japan, which is about 10 years ahead of the SET trend in the United States. You can't open a high-end audio magazine today without seeing lots of ads for very-low-powered SET amplifiers. The SET enthusiast's mantra, coined by reviewer Dick Olsher, is, "If the first watt of amplifier power doesn't sound good, why would you want 199 more of them?"

One of the early audio-amplifier triode tubes, the Western Electric 300B, was suddenly in such demand that audiophiles were paying as much as \$500 apiece for them. This shocking fact prompted Western Electric to start producing the 300B again. If you told Western Electric executives of the 1980s that in 1997 they would put the 300B back into production, they'd have thought you were crazy. (A single-ended-triode power amplifier using the 300B output tube is shown in Fig.5-3. The 300B is the bulbous tube in the middle.)

The triode is the simplest of all vacuum tubes; its glass envelope encloses just three electrical elements rather than the five elements in the more common (and modern) pentode tube. Triodes produce much less power than pentodes, but their distortion characteristics are less musically objectionable. Before the SET comeback, virtually all modern tubed amplifiers used pentodes.

In a single-ended-triode amplifier, the triode is configured so that it always amplifies the entire audio signal. That's what "single-ended" means. Virtually all other power amplifiers are "push-pull," meaning that opposing pairs of tubes (or

transistors) alternately “push” and “pull” current through the loudspeaker. (Later in this chapter we’ll look more closely at amplifier output-stage topology and class of operation.) SET proponents point to the fact that the triode amplifies the entire waveform. Moreover, SET amplifiers have no need for a circuit called a phase splitter, which makes them even simpler. Note that a single-ended tube amplifier can use more than one output tube; what makes it single-ended is that the tubes are configured in such a way that they always conduct current throughout the entire musical waveform.



Fig.5-3 A single-ended-triode monoblock amplifier based on the 300B tube (Courtesy Cary Audio)

SET circuits are further simplified because they often use very little or no *negative feedback*. Negative feedback is taking some of the amplifier’s output signal and feeding it back to the input. Such feedback lowers distortion, but many listeners have discovered that any feedback is detrimental to amplifier musicality.

On the test bench, SET amplifiers have laughably bad measured performance. They typically produce fewer than 25Wpc of output power, and have extremely high distortion—as much as 10% THD at the amplifier’s rated output. Although most SET amps use a single triode output tube, additional triodes can be

combined to produce more power. Some SET amps, however, put out only 3Wpc. In fact, there's a kind of cult around SET amplifiers that strives for lower and lower output power. Enthusiasts of these amps believe that the lower the output power, the better the sound. One SET designer told a reviewer, "If you liked my 9W amplifier, wait until you hear my 3W model." In addition to low output power and high distortion, SET amplifiers have very high output impedances, as amplifiers go: on the order of 1.5–3 ohms. This is contrasted with the 0.1 ohm output impedance of most solid-state amplifiers, and the 0.8 ohm of many push-pull tubed designs. Because a loudspeaker's impedance isn't constant with frequency, a SET amp's high output impedance reacts with the speaker's impedance variations to produce changes in frequency response. That is, the SET amplifier will have a different frequency response with every loudspeaker it drives. These variations can range from just 0.1dB with some loudspeakers that have a fairly constant impedance, to as much as 5dB with other speakers. The SET amp's sound will therefore be highly dependent on the loudspeaker with which it is matched.

Despite these technical drawbacks, my listening experience with SET amplifiers suggests that this ancient technology has many musical merits. SET amps have a certain presence and immediacy of musical communication that's hard to describe. It's as if the musicians aren't as far removed from here-and-now reality as they are with push-pull amplifiers. SET amps also have a wonderful liquidity and purity of timbre that are completely devoid of grain, hardness, and other artifacts of push-pull amplifiers. When I listen to SET amps (with the right speakers), it's as if the musicians have come alive and are playing in the listening room for me. There's a directness of musical expression that's impossible to put into words, but is immediately understood by those who have listened for themselves. You must hear a SET firsthand to know what the fuss is about; no description can convey how they sound.

When auditioning a SET amplifier, it's easy to be seduced by the midrange. That's because SET amplifiers work best in the midband, and less well at the frequency extremes of bass and treble. If the SET demo is being run for your benefit, be sure to listen to a wide variety of music, not just small-scale music or unaccompanied voice—these will accentuate the SET's strengths and hide its weaknesses. I attended an audio show at which an exhibitor of SETs played nothing but unaccompanied female vocal music for the show's entire three days.

The importance of matching a SET amplifier to the right loudspeaker cannot be overemphasized. With a low-sensitivity speaker, the SET will produce very little sound, have soft bass, and reproduce almost no dynamic range. The ideal

loudspeaker for a SET amplifier has high sensitivity (higher than 93dB/W/m), high impedance (nominal 8 ohms or higher), and no impedance dips (a minimum impedance of 6 ohms or higher). Such a speaker will produce lots of sound for a small amount of input power, and require very little current. (Remember the example earlier of Paul Klipsch's high-sensitivity speaker that produced concert-hall levels with less than 30W?) There's been a resurgence in high-sensitivity speakers that has paralleled the popularity of SET amps. Some speakers designed for SETs have sensitivities of more than 100dB, which allows them to produce satisfying listening levels with 5Wpc. SET amplifiers are often coupled with horn-loaded loudspeakers, which have extremely high sensitivity but, in my experience, often introduce unacceptable levels of coloration.

The popularity and unmistakable sound quality of SET amplifiers pose a serious dilemma: How can an amplifier that performs so poorly by every "objective" measure produce such an involving musical experience? How can 100-year-old technology eclipse, in many ways, amplifiers designed in the 21st century? What aren't we measuring in SET amplifiers that would reflect their musical magic? Why do conventional measurements fail so dismally at quantifying what's right in SET amplifiers? Do SET amplifiers sound so good because of their high distortion, or despite it? As of yet, no one has the answers to these questions.




Fig.5-4 A single-ended solid-state power amplifier (Courtesy Pass Laboratories)

Single-Ended Solid-State Amplifiers

Single-ended amplifiers aren't confined to those using ancient vacuum-tube technology. Transistors can also be configured to amplify the entire musical waveform. A solid-state, single-ended amplifier is shown in [Fig.5-4](#). Note the large heatsinks required to dissipate the additional heat produced by Class A operation.

Single-ended solid-state amplifiers have better technical performance than single-ended-triode amps, with a lower output impedance, more power, and the ability to drive a wider range of loudspeakers. They share many of the benefits of SET amps, particularly the very simple signal path, lack of crossover distortion, and greater linearity. Although single-ended solid-state amplifiers generally produce less power than their push-pull counterparts, they generally output much more power than single-ended tubed units. Nonetheless, it's a mistake to equate single-ended solid-state with single-ended tube amplifiers: there are so many other design variables that single-ended solid-state and single-ended tubed amplifiers should be considered completely different animals.

Class D (Switching) Power Amplifiers

If single-ended-triode amplifiers represent a return to fundamental technology, the Class D amplifier may represent the future of audio amplification. Class D amplifiers, also called *switching amplifiers*, have been gaining in popularity due to their small size, low weight, high efficiency, and low cost. You can hold some 250Wpc switching amplifiers in the palm of an outstretched hand. Many of them dissipate so little heat that they can be housed in an enclosed equipment cabinet—something you'd never do with a conventional amplifier, aka a linear amplifier. That's because a linear amplifier is typically about 15% efficient, meaning that it converts only about 15% of the power it draws from the wall into the electrical signal that drives the loudspeakers, dissipating the remaining 85% as heat. By contrast, a switching amplifier is as much as 90% efficient, converting just 10% of its power draw into heat. Fig.5-5 shows a switching power amplifier. This monoblock delivers 175W into 8 ohms or 335W into 4 ohms, and weighs just eight pounds.



Fig.5-5 This eight-pound switching power amplifier can deliver 175W into 8 ohms. (Courtesy NuForce)

Class D amplifiers are sometimes erroneously called “digital” amplifiers, but that appellation is reserved for a special type of Class D amplifier described later in this chapter.

At the low end of the audio market, Class D amplifiers are ubiquitous in home-theater-in-a-box systems and as integral subwoofer amplifiers. A home-theater-in-a-box may need to power six loudspeakers from a case the size of a Blu-ray player—all for a few hundred dollars. Such an amp can output perhaps 300W (50Wx6), yet run cool enough to be placed in a cabinet. In this application, the advantages of Class D are undeniable. But are Class D amplifiers suitable for high-end systems?

Before tackling that question, let's look at how a Class D amplifier works. In a conventional linear amplifier, the output transistors amplify a continuously variable analog signal—the musical waveform. The current flow through the output transistors (or tubes) is continuously variable, and is a direct analog of the musical waveform. In a Class D amplifier, the analog input signal is converted into a series of on and off pulses. These pulses are fed to the output transistors, which turn the transistors fully on or fully off. When the transistors are turned on, they conduct the DC supply voltage to the loudspeaker. When they're turned off, no voltage is connected to the loudspeaker. The audio information is contained in the durations of these on/off cycles. The train of pulses amplified by the transistors is smoothed by a filter to recover the musical waveform and remove the switching noise. Because the signal amplitude is contained in the width of the pulses, switching amplifiers are also called pulse-width-modulation (PWM) amplifiers. In fact, the Direct Stream Digital (DSD) encoding system used in the production of Super Audio Compact Discs (SACDs) is nearly identical to the PWM scheme used in Class D power amplifiers. (The graphic in [Chapter 7](#) showing how DSD's pulses represent an analog waveform is most illustrative.)

A major drawback of Class D amplifiers is their need for an extremely clean power supply. Remember that the transistors switch the DC supply directly to the loudspeaker. Consequently, any noise or ripple (a small amount of AC on the DC voltage) will appear at the loudspeaker's input terminals. Moreover, the output transistors may not turn on and off at precisely the right times, which will introduce distortion. Finally, Class D amplifiers generate large amounts of switching noise that must be filtered by large inductors and capacitors at the amplifier output. In practice, the sound quality of Class D amplifiers seems to be highly dependent on the environment, the speaker, and the speaker cables. A Class D amplifier that sounds reasonably good in one system might be unlistenable in another.

Nonetheless, some successful high-end amplifiers employ Class D technology. The field is relatively new, and manufacturers are finding ways to get good sound from Class D amplifiers. A few of the high-end Class D amps I've heard sounded excellent, suggesting that this technology may have a future in products other than car stereos and home-theaters-in-boxes.

Digital Amplifiers

A related amplifier technology uses a switching output stage, but accepts digital rather than analog input signals. These “digital” amplifiers take in the pulse-code-modulation (PCM) signal from a CD transport, music server, or other source and convert those audio data to a pulse-width-modulated signal. This PWM signal then drives the output transistors, just as in a Class D amplifier. The difference between a Class D amplifier and a digital amplifier is that the digital amplifier accepts digital data rather than an analog signal.

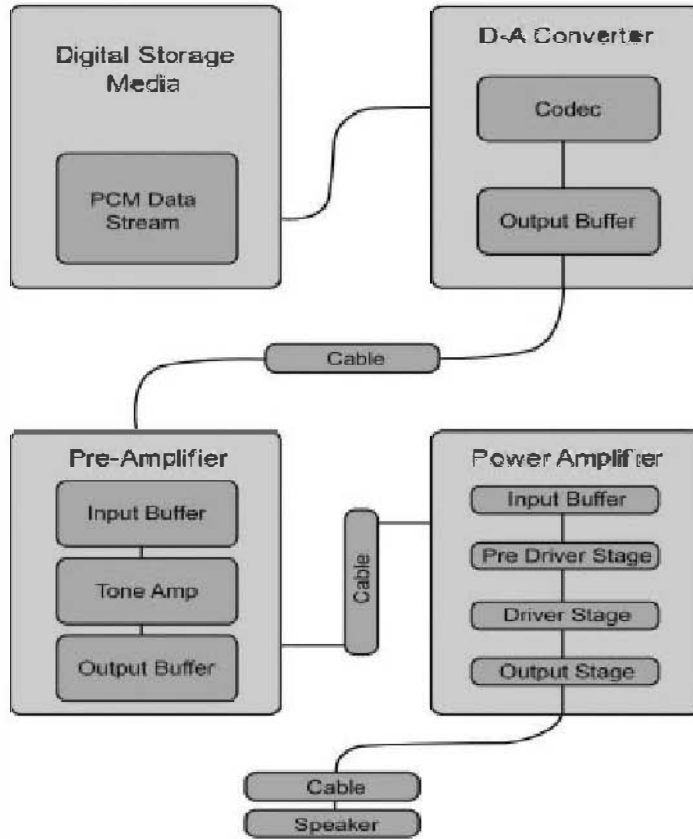
This difference might not seem that great at first glance, but consider the signal path of a conventional digital-playback chain driving a Class D power amplifier. In your music server, data read from a hard-disk drive are fed to a DAC that converts the digital data to an analog signal; the DAC’s current output is converted to a voltage by a current-to-voltage converter; the signal is low-pass filtered, then amplified in the DAC’s analog-output stage. This analog output signal travels down interconnects to a preamplifier with its several stages of amplification, volume control, and output buffer. The preamp’s output then travels down another pair of interconnects to the power amplifier, which typically employs an input stage, a driver stage, and the Class D output stage. In addition to the D/A conversion, that’s typically six or seven active amplification stages before the signal gets to the power amplifier’s output stage.

To reiterate the contrast with a true digital amplifier, PCM data are converted by DSP into the PWM signal that drives the output transistors. That’s it. There are no analog gain stages between the PCM data and your loudspeakers. The signal stays in the digital domain until the switching output stage, which, by its nature, acts as a digital-to-analog converter in concert with the output filter. The volume is adjusted in DSP. Digital amplifiers are usually not just power amplifiers, but also include a range of inputs, source selection, and volume control, effectively giving them the functional capabilities of an integrated amplifier. Fig.5-6 shows a digital amplifier, and Fig.5-7 is a block diagram comparing the signal path of this amplifier with a conventional system architecture.



Fig.5-6 A true digital amplifier accepts digital input signals for conversion to a pulse-width-modulated signal. (Courtesy NAD Electronics)

Analog Audio System



Direct Digital Audio System

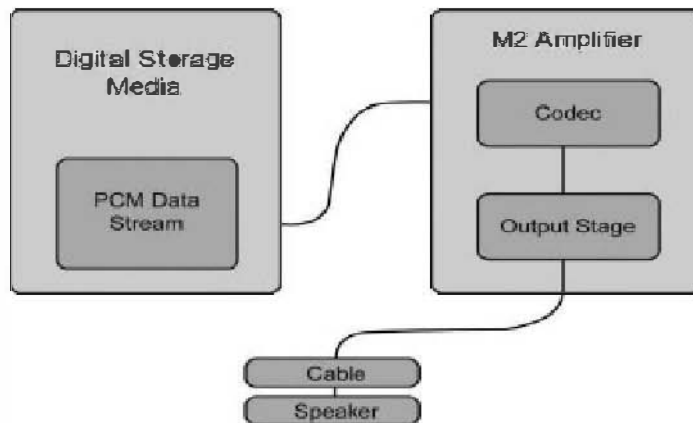


Fig.5-7 Block diagram comparing a conventional analog audio system architecture with a direct-digital amplifier architecture (Courtesy NAD Electronics)

A Unique Amplifier

In 2010, the French company Devialet introduced a radically different amplifier that may portend the future of audio amplification. The company's first product, the D Premier, featured an output stage that combined a Class A voltage amplifier with a Class D current amplifier. The idea was to merge the sonic benefits of Class A with Class D's cool operation, low power consumption, small size, and high efficiency. The company's subsequent products have built on this platform, and added many new and interesting capabilities to a category that hadn't seen much functional innovation in decades.

For example, Devialet's model 200 (Fig.5-8) performs the functions of an integrated amplifier and DAC, but with a suite of unprecedented features. Rather than think of the 200 in terms of traditional component categories, it's more useful to view it as a multifunction, general-purpose, audio hardware platform controlled by software. That hardware platform includes a 200Wpc integrated amplifier with a DAC, phono input, wireless streamer, A/D converter (with LP ripping capability), and subwoofer crossover. The 200's inputs can be configured to fit into just about any system. Don't have a turntable but have two analog line sources? No problem; the analog inputs can be configured as line inputs. Conversely, those same input jacks can become phono inputs, complete with adjustable gain and cartridge loading (impedance and capacitance). With a new technology called Speaker Active Matching (SAM), the 200's output signal can be optimized for your particular loudspeakers.



Fig.5-8 The products from Devialet defy classification by traditional component categories. (Courtesy Devialet)

To give you an idea of this product's flexibility: The Devialet 200 began life as the 170Wpc model 170. A software update turned it into the 200Wpc model 200. You can even convert the 200Wpc stereo model 200 into a 400W monoblock with the addition of a second model 200—which also doubles the number of inputs.

The 200's appearance is as radical as the technology inside. About the size and shape of a laptop computer, the chrome-plated aluminum case can be mounted flat against a wall. Only one button, an on/off switch, adorns the front panel (if you can call it that). You interact with the 200 through a square remote control with four small buttons and a huge volume knob. Alternately, you can control the 200 with an app on your tablet or mobile device.

Configuring the 200 is quite simple. A page on Devialet's website shows the 200's rear panel with the configurable components highlighted. Clicking on, for example, the RCA input jacks brings up a screen that allows you to select between line and phono; and, if phono, the cartridge gain and loading. Other phono options include mono or stereo, selectable equalization curve, and channel balance. A digital-out jack can be an analog line-out jack, with or without high- or low-pass filtering, with selectable crossover frequency and slope. The signal appearing at the binding posts can be high- or low-pass filtered, again with selectable frequency and slope. These filtering functions are ideal for systems with an active subwoofer; the signal driving the speakers is high-pass filtered, and the signal driving the sub is low-pass filtered.

Once you've configured the 200, you download the configuration file to an SD card. You then insert the SD card in the 200's rear panel, and in a few seconds the 200 has morphed into an audio product with the parameters you've specified. Because the 200 is essentially a digital platform, analog inputs, including phono signals, are digitized. You can select a sample rate of 96 or 192kHz in the A/D converter. The digitized signal, from any input, appears at a digital-out jack (RCA). You could use this function to digitize and archive a library of LPs, for example.

Digital inputs include USB, AES/EBU, TosLink, and S/PDIF coaxial. You can also wirelessly stream audio at up to 24-bit/192kHz via Devialet's AIR Universal Streamer. Wireless connection between the computer music server to the 200 has a theoretical advantage over USB, including isolating the computer's noise from the 200. The AIR app provides asynchronous connection between your computer and the 200; the signal is buffered, then processed to reduce jitter. You can also stream music files to the 200 via the 200's Ethernet connection.

Returning to Speaker Active Matching (SAM), this technology applies a correction in DSP that optimizes the output stage's current and voltage delivery for your particular speaker's impedance curve, and corrects the timing of the bass. SAM also extends the speaker's bass response through equalization, and limits the output-stage power to prevent the speaker from being overdriven. Note that SAM doesn't attempt to correct speaker- or room-induced frequency-response variations, as do full-blown DSP room-correction processors. As with all the other configuration options, SAM involves choosing your loudspeaker on the Devialet website, downloading the optimization file to an SD card, and updating the 200's software. As of this writing, Devialet has provided matching for relatively few loudspeakers, but is adding more models monthly.

Output-Stage Topology and Class of Operation

We've seen that the output stage of a power amplifier can be configured either as single-ended or push-pull. To reiterate, in a single-ended amplifier the output devices (tubes or transistors) are configured so that they always amplify the entire musical waveform. The single-ended amplifier cannot operate in any other way—that's the very definition of single-ended. In a push-pull output stage, pairs of opposing devices are arranged to work alternately—one device "pushes" current through the loudspeaker as the other device "pulls" current through the speaker. In a parallel push-pull circuit, multiple pairs of output devices are grouped together to increase the power output.

That's a description of an amplifier's output-stage topology—how the amplifying devices are configured. Now let's look at a separate but related factor: the amplifier's class of operation.

Class of operation is how a given output stage is driven. The two main classes of operation are Class A and Class B. In a Class A amplifier, the output stage (single-ended or push-pull) amplifies the entire musical waveform. In a Class B amplifier, opposing pairs of transistors (or tubes) are operated so that one transistor in the pair amplifies the positive half of the waveform as the opposing transistor amplifies the negative half of the waveform.

Class of operation is easily confused with single-ended and push-pull output-stage topologies. But there's an important distinction: A single-ended amplifier always operates in Class A, but a push-pull amplifier can be operated in Class A or Class B. It might seem as if a push-pull Class A amplifier is a contradiction in terms, but it's not. In a push-pull Class A amplifier, opposing pairs of transistors are driven in such a way that current flows through both transistors throughout the entire musical waveform, and all the output devices participate in amplifying through the full cycle of the audio signal. One device pulls current through the loudspeaker while the other pushes current, but both are always turned on and conducting current.

By contrast, in a Class B push-pull output stage, one transistor amplifies the signal during the positive-going portion of the signal, and the other amplifies the negative-going half of the signal. When one is working, the other is resting (and getting some needed cooling).

The two factors are the output-stage topology (how the output devices are configured) and the class of operation (how those output devices are driven). The following four statements may clarify the matter:

- Single-ended topology always operates in Class A
- Class A operation can be implemented in a single-ended or push-pull topology
- Push-pull topology can be operated in Class A or Class B
- Class B operation is always implemented by a push-pull topology

Let's move on from these distinctions to the more practical aspects of amplifier class of operation.

Most power amplifiers are categorized as Class A/B because they operate in Class A at very low power outputs, then default to Class B at higher power outputs. A 100Wpc amplifier may put out 5W of Class A power, and then switch to Class B above that level. Even the heftiest Class A/B amplifiers can put out only a small portion of their rated power in Class A—a typical value is 1 or 2%. Although this may not sound like much, an amplifier driving high-sensitivity speakers at low listening levels may be delivering only a couple of watts.

How much of an amplifier's power output is Class A is determined by the amount of bias applied to the output transistors. Bias is a DC current that flows through the output stage at idle (no audio signal). The higher the bias, the more current flows through the transistors when no signal is present. More bias results in more Class A output power. Class B operation has no bias current; Class A/B has moderate bias current; and Class A has high bias current. The designer can keep increasing the bias in a push-pull output stage until all of its output-power capability is delivered in Class A. The amplifier's power-output rating would be the point at which the amplifier leaves Class A and begins operating in Class B. The factors limiting increasing the bias current are the ability of the transistors to handle the greatly increased current flow through them, of the power supply to keep up with the transistors' current demands, and of the heatsinking to dissipate the considerable heat caused by the high bias current.

To give you an idea of the demands placed on an amplifier's output stage, power supply, and heatsinking by an output stage biased for Class A operation, let's compare two integrated amplifiers from Pass Labs that have identical output stages but are biased completely differently: the INT-150, 150Wpc in Class A/B; and the INT-30A, 30Wpc in pure Class A. The INT-150's push-pull output stage is biased so that it produces 10Wpc of Class A power before switching to Class B

to deliver its full rated output of 150Wpc. The amplifier can double its output power, to 300Wpc into 4 ohms. The INT-30A is exactly the same amplifier, employing the same power supply, an identical number and type of output transistor, and the same heatsinks as the INT-150, but is rated at just 30Wpc. The difference is that the INT-30A's 30Wpc are pure Class A watts. The amplifier, biased into Class A, delivers just one-fifth the power output of its Class A/B counterpart.

As you can see, Class A operation is enormously inefficient. A Class A amplifier converts nearly all the power it draws from the wall outlet into heat, and consumes just as much power at idle as when it is operating at its maximum output power. Moreover, a Class A amplifier is much more expensive to build on a "watts per dollar" basis than its Class A/B counterpart. So why would designers go to the trouble and expense of creating Class A amplifiers, and why would consumers pay such a huge premium for "Class A watts" over "Class A/B watts"?

Class A has many theoretical and practical advantages. For starters, Class B and A/B amplifiers suffer from crossover distortion, a discontinuity at the zero-crossing point where one transistor in each opposed pair "hands off" the signal to its partner. A waveform discontinuity can occur at this transition, and is lessened as bias current is introduced and increased in value. Crossover distortion can't occur in Class A operation because each transistor amplifies the entire audio waveform, not just half of it. Second, the large thermal hardware capacity required by Class A has the advantage of keeping the output transistors more thermally stable (i.e., at a more constant temperature). This makes their operating characteristics more uniform, and less subject to changes resulting from the signal characteristics the transistors are amplifying. For example, if the transistors have just delivered a surge of current to the loudspeakers, they won't behave differently—and thus sound different—immediately afterward because they are momentarily hotter. Third, increasing the bias current so that an amplifier produces more Class A power not only reduces harmonic distortion but, more important, changes the nature of the harmonic distortion. As the bias is increased, the upper-order harmonics (everything above the third harmonic), which are the most sonically harmful, are most dramatically reduced in amplitude, leaving the predominant distortion component the more sonically benign second and third harmonics.

Class A power amplifiers can sound extremely good, with a sweetness and liquidity that set them apart from Class A/B amps. In my experience, Class A amplifiers have many of the virtues of a tubed amplifier but without the tube amp's technical limitations. This isn't to say that a Class A amplifier mimics the

sonic character of tubes, but rather that Class A avoids many of the characteristic nonlinear distortions of Class A/B solid-state amplifiers.

How a Power Amplifier Works

It's intuitive to conceptualize a power amplifier as a device that simply makes the input signal bigger. A more accurate visualization is of a device that pulls 120V, 60Hz alternating current (AC) from your wall outlet through the amplifier's power transformer, converts the AC into direct current (DC), stores that energy in large capacitors (the can-shaped objects), then allows the tiny audio signal at the amplifier input to modulate the stored energy as electrical current that is then driven through the speakers by the power amplifier's output transistors or tubes.

This way of thinking about a power amplifier produces the realization that a power amplifier's power supply (the section that converts the AC from the wall outlet to the DC that supplies the output transistors) is actually in the audio signal path. The current that ultimately drives the cones in your speakers back and forth comes from the wall outlet via the amplifier's power supply. That's one reason a power amplifier's power supply is so important—and why good power amplifiers tend to have large, heavy power supplies.

At its most basic level, a power amplifier consists of a power supply, an input stage, a driver stage, and an output stage. The input and driver stages prepare the signal for the output stage, which is where transistors or tubes drive electrical current through the loudspeaker. Let's look at each of these sections in more detail.

The Power Supply

The *power supply* converts 60Hz, 120V AC from the wall outlet into the DC the audio circuits need. Power-amplifier power supplies differ greatly from those for other audio components. (A close reading of [Appendix B](#) will be helpful in understanding the rest of this explanation.) Power-amplifier supplies must be able to provide large amounts of current, operate at much higher voltages, and have much larger storage capacities. Parts sizes and ratings are thus much higher in a power amplifier than in the power supply of a preamp or disc player.

The first element of the power supply is the *power transformer*, which steps down the 120V AC line voltage to a lower value, perhaps 80V. You can get a pretty good idea of how powerful the amplifier is by the power transformer's size and rating. Transformers are rated in "VA" or "kVA" (x 1000); i.e., voltage (V) multiplied by the maximum current in amperes (A) they can supply. A transformer with a VA rating of 1kVA can deliver 8.34 amperes of current (1000 divided by the AC line voltage of 120V) to the amplifier's circuits (assuming the transformer isn't stepping down the voltage). Because most power transformers step down the voltage to 80V, a 1kVA transformer could deliver 12A.

Note that the amplifier's overall continuous output current for both channels combined can never exceed the power transformer's current rating. The amplifier can, however, deliver short bursts of higher current (such as on musical transients like bass-drum whacks) because of the energy stored in the capacitors—provided the output transistors can handle the current. One very large power amplifier's transformer, with a VA rating of 4.5kVA, weighs 85 of the amplifier's total of 147 pounds. A power amplifier of more moderate size may have a power transformer with a rating of 300VA.

This stepped-down AC voltage from the power transformer is then rectified (turned into DC) by a bridge rectifier and filtered by very large filter capacitors. The filter capacitors not only store energy, they smooth the DC by removing any trace of ripple (a little bit of AC appearing in the DC supply). The filter capacitors are also called reservoir capacitors because they store the large amounts of electrical energy needed to run the amplifier. This reservoir of electrical energy is needed because the power supply can't pull current from the wall outlet fast enough to deliver instantaneous surges of high current to the speakers. The ability to store a charge is called capacitance, and its unit of measurement is the *farad*, named after the English chemist and physicist Michael Faraday (1791–1867). Some capacitors are the size of soda cans, and can store huge amounts of

electrical energy, typically between 30,000 and 500,000 microfarads (μF). If you turn off such a power amplifier when music is playing, you'll hear music for a few seconds, or nearly a minute. One would intuitively think that the reservoir capacitors are refilled continuously by the 120V, 60Hz AC from the wall. In reality, the amplifier draws current from the wall to refill the capacitors for only a few milliseconds, at the sinewave's positive and negative peaks. This fact has important implications for how we connect the amplifier to the AC wall outlet, as we'll see in [Chapter 15](#)'s section on AC power cords and conditioners.

This high-energy supply is the power source for the amplifier's output stage; the electrons in the filter capacitors are pulled out by the output transistors and delivered as current flow to the loudspeaker. The small audio signal at the power-amplifier input modulates the release of the energy stored in the power supply; this energy release is the amplifier's output signal, which drives the speakers.

The ideal power supply would maintain its voltage no matter how much current the output transistors demanded. In practice, the output-supply voltage is pulled down when the output transistors suddenly draw current, as happens during a musical peak. This drop in power-supply output voltage is called droop. If your lights have ever dimmed momentarily when the refrigerator motor turned on, you've seen the results of power-supply droop. The refrigerator motor draws a large amount of current when it starts, pulling down the voltage available to power the lights. A power supply that resists drooping is called a stiff supply.

The fuses on the backs of many power amplifiers are wired in series with the output stage's power supply. That is, the current drawn from the power supply by the output transistors must flow through the fuses. Because the power-supply voltage is called the voltage rail, these fuses are called *rail fuses*. If the output stage draws too much current, the fuse blows. If a fuse blows once, replace it with one of identical rating. Under no circumstances should you replace a fuse with one of a higher rating. The fuse is there for a reason: to protect the output stage and the power supply. Defeating the fuse, or replacing a lower-value fuse with a higher-value one, can cause catastrophic damage to the amplifier. Moreover, there's a reason if the fuse blows a second time: something is wrong in the amplifier. Fuses often blow when an output transistor goes bad; have the amplifier serviced by a qualified technician.

This description is of the conventional sort of linear power supply used in most amplifiers. An alternative is the switched-mode power supply, in which the transistors that conduct current to the load turn on and off rather than stay on constantly, as in a linear supply. The advantage is smaller size, lower weight, less heat, and greater efficiency.

Another term related to an amplifier's power supply is power-factor correction. Without power-factor correction, the reactance of the power supply to the AC line causes the voltage and current to be slightly out of sync. This phenomenon reduces the amount of real power the supply can deliver to the load. A power supply that has been power-factor-corrected compensates for this problem, resulting in a more efficient power supply. My listening experience with amplifiers employing supplies with power-factor correction suggests that they exhibit exceptional bass solidity, dynamics, and authority.

An amplifier's power supply also generates lower voltages to power the less power-hungry circuits in the amplifier, such as the input and driver stages (described next).

Input and Driver Stages

The audio input signal driving the power amplifier is first buffered and amplified by the power amplifier's input stage. The input stage has a high input impedance and a low output impedance, making the amplifier easy to drive by the preamplifier. It is here, at the input stage, that the amplifier's input impedance is determined. Specified in ohms, input impedance is the resistance to current flow of the preamplifier's drive signal. An input impedance of 47k ohms is fairly standard for an unbalanced input on an RCA jack.

The input stage is followed by the driver stage, the circuit that amplifies the signal and drives the final output stage. The driver stage produces a higher-voltage signal that the output transistors need. In Class A/B amplifiers, the driver stage can act as a phase splitter to create an identical copy of the signal (with polarity inverted) for presentation to the push-pull output stage.

Many amplifiers have *coupling capacitors* between these stages. A coupling capacitor, also called a *DC blocking capacitor*, is a barrier to direct current (DC), but passes the alternating current (AC) audio signal to the next stage. But the fact that these capacitors can degrade the amplifier's musical performance leads some designers to omit them from the signal path. An amplifier without coupling capacitors is called *direct-coupled*. Direct-coupled amplifiers can present a danger: any DC at the input will be amplified and sent to your speakers. Because even small amounts of DC can destroy a loudspeaker, some amplifiers have a circuit called a DC servo to prevent this DC from appearing at the output terminals. A DC servo monitors the DC level and adjusts the circuitry for minimum DC. You can easily measure how much DC appears at a power amplifier's outputs: just put a DC voltmeter across the speaker terminals. Values below 200mV are typical.

Most high-quality power amplifiers use *discrete* circuitry in the input and driver stages. A discrete circuit is made from individual transistors rather than integrated-circuit operational amplifiers (op-amps). Discrete circuits are generally better than op-amp-based designs. *Field Effect Transistors* (FETs) are often used in the input stages of power amplifiers for their intrinsically high input impedance. Their low power capability, however, limits their use to the low-level signals of the input stage. Input and driver stages are also often operated in Class A, meaning that the transistor amplifies the entire musical waveform.

Nearly all power amplifiers use feedback, a technique in which part of the output signal is returned to the input. This makes the circuit more linear (the

input and output signals have greater similarity) and lowers distortion. Too much feedback, however, results in poor sound. Many power amplifiers are touted as “low-feedback” designs. These amplifiers will probably have higher distortion figures, but may sound better as a result of lower feedback.

Two types of feedback are local and global. Local feedback is feedback used within only one stage of an amplifier; global feedback is feedback from the final output stage back to the input stage. Finally, the input and driver stages are almost always run in Class A, meaning that one transistor amplifies the entire audio signal.

Output Stage

In discussions of power amplifiers, the most common topic is the *output stage*. The output stage comprises many transistors, usually bipolar types, and, less frequently, *MOSFETs*, another type of transistor. These transistors are the workhorses of a power amplifier—the current that makes the loudspeaker cones move is driven by the output transistors. Consequently, adding more transistors to the output stage increases the amplifier’s output-current capability. More transistors, however, require a bigger power supply. In addition, these output transistors produce lots of heat that must be dissipated by *heatsinks*, the large black fins that protrude from the sides of many power amplifiers.

When the output stage is asked to deliver more power than it is capable of, the top and bottom of the waveform are flattened, or *clipped*. When a power amplifier goes into clipping, the output is severely and audibly distorted. If clipping occurs only briefly on a few musical peaks, we may not notice it. An amplifier driven into hard clipping will, however, quickly destroy loudspeakers. Loudspeakers can handle surprising amounts of power, provided that the power amplifier driving them isn’t itself driven into clipping.

Some power amplifiers use MOSFET output devices instead of bipolar transistors. MOSFET stands for Metal Oxide Semiconductor Field Effect Transistor. These devices operate differently: they are turned on (biased) by voltage rather than by current, as are bipolar devices. In this respect, they operate like vacuum tubes. MOSFETs have other technical attributes, and MOSFET amplifiers often have a characteristic sound, particularly a smoothness in the treble. They also exhibit what’s been called “MOSFET mist,” a tendency to have a less than crystal-clear soundstage and a slightly unfocused treble.

How a Tubed Power Amplifier Works

Though the basic functions of tubed and solid-state power amplifiers are almost identical, instead of using transistors in the input, driver, and output stages, tubed power amplifiers use vacuum tubes. A tubed power amplifier may use one tube in the input stage, another in the driver stage, and as many as eight in a push-pull-configured output stage. Hybrid power amplifiers use tubes in the input and driver stages, and solid-state devices in the output stage.

Fig.5-9 shows a stereo tubed power amplifier. Other than the tubes themselves, the most obvious difference between tubed and solid-state power amplifiers is that a tubed amp has one output transformer per channel, clearly visible at the amplifier's rear—a feature not found in solid-state amplifiers. The output transformer functions as an impedance-matching device between the amplifier's output stage and the loudspeaker. The tubed amplifier's output stage thus “sees” the transformer's high input impedance rather than the loudspeaker's low input impedance. The transformer also cancels harmonic distortion by mixing the plate currents generated by both halves of the push-pull output stage. Finally, the output transformer prevents DC from appearing at the amplifier output terminals.

The *output transformer* must be large enough not to saturate at low frequencies. Saturation is when the transformer has reached its operational limits. If saturation occurs, the bass sounds soft and “woolly.” Some tubed power amplifiers have selectable taps on their transformers for loudspeakers of different impedances. These outputs are additional transformer windings called secondary windings, and provide the optimal impedance match between transformer and loudspeaker.



Fig.5-9 A tubed stereo power amplifier has a power transformer and two output transformers, seen as the large cubes at the rear. (Courtesy Mystère)

A few tubed amplifiers have no output transformers and thus are called *output-transformerless* (OTL). OTLs generally aren't very powerful, and don't work well with low-impedance loudspeakers. They're huge, have lots of output tubes, and are idiosyncratic in loudspeaker matching. The OTL amps I've heard have mediocre bass and not much in the way of dynamic impact, but spectacular soundstaging and beautiful rendering of timbre. OTLs aren't for everyone, but they can sound magical with the right loudspeakers.

Because a tube's operating characteristics change with age, the biases of all output tubes need to be adjusted periodically. *Bias* is a small amount of current, typically between 20 and 70mA, applied to the tube's grid to turn on the tube. Its transistor counterpart is the bias current we talked about in solid-state output stages under "Output-Stage Topology and Class of Operation." A tubed amplifier will often have test points and trimmer potentiometers, called trim pots—tiny variable resistors adjustable with a small screwdriver—readily accessible on the front panel or chassis for adjustment of bias.

To set an output tube's bias, put a DC voltmeter probe on the test point, ground the other test lead, and turn the trim pot until the voltage reading matches the manufacturer's bias specification. Bias is especially important in amplifiers having many output tubes; we want the tubes to operate identically. Not only do tubes age differently, but, due to variations in manufacturing, they need different amounts of bias even when new. When an output tube approaches

the end of its useful life, it may not reach the specified bias voltage even when its bias trim pot is turned all the way up. When that happens, it's time for a new tube. Bias should be checked about once a month, or whenever changing tubes. Some tubed amplifiers have an automatic bias-calibration feature that sets each tube's bias automatically. Others offer a front-panel meter or LED to indicate the bias level. The latter feature, by obviating the need for a voltmeter, makes owning a tubed amplifier more convenient. In amplifiers with an LED indicator, the LED will illuminate when the bias is set within the correct operating range.

When you buy replacement output tubes, you can specify matched pairs or quartets of tubes. These are tubes whose electrical characteristics have been matched at the factory or by the tube reseller for better combined performance. An amplifier with a single push-pull pair per channel will benefit from matched pairs, and an amplifier with four output tubes per channel needs matched quartets. Matched output tubes ensure that each tube operates very similarly, and that one isn't working harder than another. An analogy is synchronizing the carburetors of a multicylinder motorcycle engine. Expect to pay a premium for matched tubes, however.

A tubed amplifier is often referred to by the type of tube used in its output stage: triode or pentode. Triode power amplifiers generally sound best but aren't very powerful, and are much more expensive per watt of output power than other tubed designs. Triode-based amplifiers can be of the push-pull variety, or the single-ended type described earlier.

A pentode output tube has five components within the glass envelope, which greatly increases the tube's output power and efficiency. Some tubed amplifiers let you switch between pentode and triode operation using the same output tube. You may choose the purer triode mode for chamber music at low volume, and pentode for louder music that requires more output power. The switchable option also lets you select the best output mode for your loudspeakers. High-sensitivity speakers may work just fine when driven by the amplifier in triode mode; lower-sensitivity models may require the higher output power provided by pentode operation.

Another type of tubed output stage, called *ultralinear*, was developed in 1951 by David Hafler and Herbert Keroes. Ultralinear operation is a method of connecting a pentode so that it operates more like a triode. Specifically, some of the plate voltage is returned to the pentode's screen grid. Ultralinear operation was the dominant circuit topology in tubed amplifiers for decades because it offered the output power of pentodes with the distortion characteristics of triodes.

Some tubed amplifiers, particularly the single-ended variety, let you vary the amount of feedback. You can thus dial in the sound you want, and decide which musical trade-offs are acceptable. Reducing the amount of feedback often results in purer timbre, but at the expense of bass tautness.

Note that tubed amplifiers don't increase their power outputs as the load impedance drops, as do solid-state designs. The output transformer acts as a current source rather than as a voltage source, decreasing the available output voltage when the load impedance drops. This is one reason tubed amplifiers have less tight, less extended bass than solid-state power amps.

Finally, the power supplies of tubed power amplifiers are different from their solid-state counterparts. The voltage at a tubed amp's output-stage supply rail is many hundreds of volts, not the $\pm 60\text{V}$ to $\pm 100\text{V}$ of a solid-state amp's power supplies. Because of this very high voltage, the reservoir capacitors generally store vastly more energy than the reservoir caps in a solid-state amp. This so-called "B+" voltage is very high—the capacitors can deliver lots of current through your body, and can even kill you. Energy storage is measured in joules, and is a function of the capacitance in farads (a unit of stored charge) multiplied by the voltage squared. Given the same capacitance (same number of microfarads) but ten times the voltage, the power supplies of a tubed amplifier can store much more energy than those of a solid-state amp. This factor may account for tube amplifiers' subjectively greater loudness for the same rated power compared to a solid-state amplifier.

Tubed amplifiers must also supply the tube filaments (the part that glows) with a low DC voltage. In low-signal-level tubes, the number beginning the tube's designation (6DJ8, 12AX7, etc.) indicates the filament voltage, usually 6V or 12V. Filaments are also called by the more descriptive name "heaters."

From a design standpoint, tubed power amplifiers have several inherent advantages. First, they're simple. Comparing the schematics of solid-state and tubed amplifiers reveals many more parts in the transistor design, most of them in the signal path. Fewer parts in the signal path generally results in better sound. When overdriven, tubes soft-clip, meaning that the waveform gently rounds off rather than flattening immediately. Moreover, when clipping, a tubed amplifier produces primarily second- and third-harmonic distortion, which is sonically fairly benign. A transistor amplifier driven into clipping immediately produces a whole series of very objectionable upper-order odd harmonics (fifth, seventh, ninth). In fact, 10% second-harmonic distortion is less annoying than 0.5% seventh-harmonic distortion.

6 Loudspeakers

Introduction

Of all the components in your audio system, the loudspeaker's job is by far the most difficult. The loudspeaker is expected to reproduce the sound of a pipe organ, the human voice, and a violin through the same electromechanical device—all at the same level of believability, and all at the same time: the tonal range of virtually every instrument in the orchestra is reproduced from a relatively tiny box. This frequency span of ten octaves represents a sound-wavelength difference of 60 feet in the bass to about half an inch in the treble.

It's no wonder that loudspeaker designers spend their lives battling the laws of physics to produce musical and practical loudspeakers. Unlike other high-end designers who create a variety of products, the loudspeaker designer is singular in focus, dedicated in intent, and deeply committed to the unique blend of science and art that is loudspeaker design.

Although even the best loudspeakers can't convince us that we're hearing live music, they nonetheless are miraculous in what they can do. Think about this: a pair of loudspeakers converts two two-dimensional electrical signals into a three-dimensional "soundspace" spread out before the listener. Instruments seem to exist as objects in space: we hear the violin here, the brass over there, and the percussion behind the other instruments. A vocalist appears as a palpable, tangible image precisely between the two speakers. The front of the listening room seems to disappear, replaced by the music. It's so easy to close your eyes and be transported into the musical event.

To achieve this experience in your home, however, you must carefully choose the best loudspeakers from among the literally thousands of models on the market. As we'll see, choosing loudspeakers is a challenging job.

How to Choose a Loudspeaker

The world abounds in poor-quality, even dreadful, loudspeakers. What's more, some very bad loudspeakers are expensive, while superlative models may sell for a fraction of an inferior model's price. There is sometimes little relationship between price and musical performance.

This situation offers the loudspeaker shopper both promise and peril. The promise is of finding an excellent loudspeaker for a reasonable price. The peril is of sorting through mediocre models to find the rare gems that offer either high absolute performance, or sound quality far above what their price would indicate.

This is where reviews come in handy. Reviewers who write for audio magazines hear lots of loudspeakers (at dealers, trade shows, and consumer shows), but review only those that sound promising. This weeds out the vast majority of underachievers. Of the loudspeakers that are reviewed, some are found to be unacceptably flawed, others are good for the money, while a select few are star overachievers that clearly outperform their similarly priced rivals.

The place to start loudspeaker shopping, therefore, is in the pages of a reputable magazine with high standards for what constitutes good loudspeaker performance. Be wary of magazines that end every review with a "competent for the money" recommendation. Not all loudspeakers are good; therefore, not all reviews should be positive. The tone of the reviews—positive or negative—should reflect the wide variation in performance and value found in the marketplace.

After you've read lots of loudspeaker reviews, make up your short list of products to audition from the *crème de la crème*. There are several criteria to apply in making this short list to ensure that you get the best loudspeaker for your individual needs. As you apply each criterion described, the list of candidate loudspeakers will get shorter and shorter, thus easing your decision-making process. If you find yourself with too few choices at the end of the process, go back and revise your criteria. For example, if you find a loudspeaker that's perfect in all ways but size, you may want to find the extra space in your living room. Similarly, an ideal loudspeaker costing a little more than you planned to spend may suggest a budget revision. As you go through this selection process, remember that the perfect loudspeaker for you is probably out there. Be selective and have high standards. You'll be rewarded by a much higher level of musical performance than you thought you could afford.

1) Size, Appearance, and Integration in the Home

After you've designated a place for your loudspeakers, determine the optimal speaker size for your room—the urban apartment dweller will likely have tighter size constraints than the suburban audiophile. Some listeners will want the speakers to discreetly blend into the room; others will make the hi-fi system the room's center of activity and won't mind large, imposing loudspeakers. When choosing a place for your speakers, keep in mind that their placement is a crucial factor in how your system will sound. (Chapter 14 includes an in-depth treatment of speaker positioning.)

The loudspeaker's appearance is also a factor to consider. An inexpensive, vinyl-covered box would be out of place in an elegantly furnished home. Many high-end speakers are finished in beautiful cabinetry or automotive paint finishes that will complement any decor. This level of finish can, however, add greatly to the loudspeaker's price.

If you don't have room for full-range, floorstanding speakers, consider a separate *subwoofer/satellite system*. This is a speaker system that puts the bass-reproducing driver in an enclosure you can put nearly anywhere, and the midrange- and treble-reproducing elements in a small, unobtrusive cabinet. You'll still get a full sound, but without the visual domination of your living room that often goes with floorstanding speakers. Moreover, the satellite speakers' small cabinets often help them achieve great soundstaging.

Although the term “bookshelf” is often applied to small speakers, you can't get optimal performance from a speaker mounted in a bookshelf. Small speakers need to be mounted on stands and placed out in the room. Small loudspeakers mounted on stands, sometimes called minimonitors, often provide terrific imaging, great clarity in the midband and treble, and can easily “disappear” into the music. On the down side, small speakers used without a subwoofer have restricted dynamics, limited bass extension, and won't play as loudly as their floorstanding counterparts.

2) Match the Loudspeaker to Your Electronics

The loudspeaker should be matched to the rest of your system, both electrically and musically. A loudspeaker that may work well in one system may not be ideal for another system—or another listener.

Let's start with the loudspeaker's electrical characteristics. The power amplifier and loudspeaker should be thought of as an interactive combination:

the power amplifier will behave differently when driving different loudspeakers. Consequently, the loudspeaker should be chosen for the amplifier that will drive it.

The first electrical consideration is a loudspeaker's *sensitivity*—*how much sound it will produce for a given amount of amplifier power*. Loudspeakers are rated for sensitivity by measuring their sound-pressure level (SPL) from one meter away while they are being fed one watt (1W) of power. For example, a sensitivity specification of “88dB/1W/1m” indicates that this particular loudspeaker will produce a sound-pressure level of 88dB when driven with an input power of 1W, measured at a distance of 1m. High-end loudspeakers vary in sensitivity between 80dB/1W/1m and 106dB/1W/1m.

A loudspeaker's sensitivity is a significant factor in determining how well it will work with a given power-amplifier output wattage. To produce a loud sound (100dB), a loudspeaker rated at 80dB sensitivity would require 100W. A loudspeaker with a sensitivity of 95dB would require only 3W to produce the same sound-pressure level. Each 3dB decrease in sensitivity requires double the amplifier power to produce the same SPL. (This is discussed in greater technical detail in [Chapter 5](#), “Power Amplifiers.”)

Another electrical factor to consider is the loudspeaker's load *impedance*. This is the electrical resistance the power amp meets when driving the speaker. The lower the speaker's impedance, the more demand it places on the power amp. If you choose low-impedance speakers, be certain the power amp will drive them adequately. (See [Chapter 5](#) for a full technical discussion of loudspeaker impedance as it relates to amplifier power.)

On a musical level, you should select as sonically neutral a speaker as possible. If you have a bright-sounding CD player or power amp, it's a mistake to buy a loudspeaker that sounds soft or dull in the treble to compensate. Instead, change your CD player or amplifier.

Another mistake is to drive high-quality speakers with poor amplification or source components. The high-quality loudspeakers will resolve much more information than lesser speakers—including imperfections in the electronics and source components. All too many audiophiles drive great speakers with mediocre source components and never realize their loudspeakers' potential. Match the loudspeakers' quality to that of the rest of your system. (Use the guidelines in [Chapter 2](#) to set a loudspeaker budget within the context of the cost of your entire system.)

3) Musical Preferences and Listening Habits

If the perfect loudspeaker existed, it would work equally well for chamber music and heavy metal. But because the perfect loudspeaker remains a mythical beast, musical preferences must play a part in choosing a loudspeaker. If you listen mostly to small-scale classical music, choral works, or classical guitar, a minimonitor would probably be your best choice.

Conversely, rock listeners need the dynamics, low-frequency extension, and bass power of a large full-range system. Different loudspeakers have strengths and weaknesses in different areas; by matching the speaker to your listening tastes, you'll get the best performance in the areas that matter most to you.

Other Guidelines in Choosing Loudspeakers

In addition to these specific recommendations, there are some general guidelines you should follow in order to get the most loudspeaker for your money.

First, buy from a specialty audio retailer who can properly demonstrate the loudspeaker, advise you on system matching, and tell you the pros and cons of each candidate. Many high-end audio dealers will let you try the loudspeaker in your home with your own electronics and music before you buy.

Take advantage of the dealer's knowledge—and reward him with the sale. It's not only unfair to the dealer to use his or her expensive showroom and knowledgeable salespeople to find out which product to buy, and then look for the loudspeaker elsewhere at a lower price; it also prevents you from establishing a mutually beneficial relationship with the dealer.

In general, loudspeakers made by companies that make *only* loudspeakers are better than those from companies that also make a full line of electronics. Loudspeaker design may be an afterthought to the electronics manufacturer—something to fill out the line. Conversely, many high-end speaker companies have an almost obsessive dedication to the art of loudspeaker design. Their products' superior performance often reflects this commitment. There are, however, a few companies that produce a full line of products, including loudspeakers, that work well with each other.

Don't buy a loudspeaker based on technical claims. Some products claiming superiority in one aspect of their performance may overlook other, more important aspects. Loudspeaker design requires a balanced approach, not reliance on some new “wonder” technology that may have been invented by the loudspeaker manufacturer's marketing department. Forget about the technical hype and listen to how the loudspeaker reproduces music. You'll hear whether or not the loudspeaker is any good.

Don't base your loudspeaker purchases on brand loyalty or longevity. Many well-known and respected names in loudspeaker design of 20 years ago are no longer competitive. Such a company may still produce loudspeakers, but its recent products' inferior performance only throws into relief the extent of the manufacturer's decline. The brands the general public thinks represent the state of the art are actually among the worst-sounding loudspeakers available. These companies were either bought by multinational business conglomerates who

didn't care about quality and just wanted to exploit the brand name, or the company has forsaken high performance for mass-market sales.

The general public also believes that the larger the loudspeaker and the more drivers it has, the better it is. But given the same retail price, there is often an inverse relationship between size/driver count and sonic performance. A good two-way loudspeaker—one that splits the audioband into two parts for reproduction by a woofer and a tweeter—with a 6" midrange/woofer and a tweeter in a small cabinet is likely to be vastly better than a similarly priced four-way in a large, floorstanding enclosure. Two high-quality drivers are much better than four mediocre ones. Further, the larger the cabinet, the more difficult and expensive it is to make it free from vibrations that degrade the sound. The four-way speaker's more extensive crossover will require more parts; the two-way can use just a few higher-quality crossover parts. The large loudspeaker will probably be unlistenable; the small two-way may be superbly musical.

If both of these loudspeakers were shown in a catalog and offered at the same price, however, the large, inferior system would outsell the high-quality two-way by at least ten times. The perceived value of more hardware for the same money is much higher.

The bottom line: You can't tell anything about a loudspeaker until you listen to it. In the next section, we'll examine common problems in loudspeakers and how to choose one that provides the highest level of musical performance.

Finding the Right Loudspeaker—Before You Buy

You've done your homework, read reviews, and narrowed down your list of candidate loudspeakers based on the criteria described earlier—you know what you want. Now it's time to go out and listen. This is a crucial part of shopping for a loudspeaker, and one that should be approached carefully. Rather than buying a pair of speakers on your first visit to a dealer, consider this initial audition to be simply the next step. Don't be in a hurry to buy the first loudspeaker you like. Even if it sounds very good to you, you won't know *how* good it is until you've auditioned several models.

Audition the loudspeaker with a wide range of familiar recordings of your own choosing. Remember that a dealer's strategic selection of music can highlight a loudspeaker's best qualities and conceal its weaknesses—after all, his job is to present his products in the best light. Further, auditioning with only audiophile-quality recordings won't tell you much about how the loudspeaker will perform with the music you'll be playing at home, most of which was likely not recorded to high audiophile standards. Still, audiophile recordings are excellent for discovering specific performance aspects of a loudspeaker. The music selected for auditioning should therefore be a combination of your favorite music, and diagnostic recordings chosen to reveal different aspects of the loudspeaker's performance. When listening to your favorite music, forget about specific sonic characteristics and pay attention to how much you're enjoying the sound. Shift into the analytical mode only when playing the diagnostic recordings. Characterize the sound quality according to the sonic criteria described in [Chapter 3](#) and later in this chapter.

Visit the dealer when business is slow so you can spend at least an hour with the speaker. Some loudspeakers are appealing at first, then lose their luster as their flaws begin to emerge over time. The time to lose patience with the speakers is in the dealer's showroom, not a week after you've bought them. And don't try to audition more than two sets of speakers in a single dealer visit. If you must choose between three models, select between the first two on one visit, then return to compare the winner of the first audition with the third contender. You should listen to each candidate as long as you want (within reason), to be sure you're making the right purchasing decision.

Some loudspeakers have different tonal balances at different listening heights. Be sure to audition the speaker at the same listening height as your listening chair at home. A typical listening height is 36", measured from the floor to your ears.

Further, some loudspeakers with first-order crossovers sound different if you sit too close to them. When in the showroom, move back and forth a few feet to be certain the loudspeaker will sound the same as it should at your home listening distance.

Make sure the speakers are driven by electronics and source components of comparable quality to your components. It's easy to become infatuated with a delicious sound in a dealer's showroom, only to be disappointed when you connect the speakers to less good electronics. Ideally, you should drive the speakers under audition with the same level of power amp as you have at home, or as you intend to buy with the loudspeakers.

Of course, the best way to audition loudspeakers is in your own home—you're under no pressure, you can listen for as long as you like, and you can hear how the speaker performs with your electronics and in your listening room. Home audition removes much of the guesswork from choosing a speaker. But because it's impractical to take every contender home, and because many dealers will not allow this, save your home auditioning for only those speakers you are seriously considering.

What to Listen For

There are several common flaws in speaker performance that you should listen for. Though some of these flaws are unavoidable in the lower price ranges, a speaker exhibiting too many of them should be quickly passed over.

Listen for thick, slow, tubby bass. One of the most annoying characteristics of poor speakers is colored, peaky, and pitchless bass. You should hear distinct pitches in bass notes, not a low-frequency, “one-note” growling under the music. Male speaking voice is a good test for upper-bass colorations; it shouldn’t have an excessive or unnatural chesty sound.

Certain bass notes shouldn’t sound louder than others. Listen to solo piano with descending or ascending lines played evenly in the instrument’s left-hand, or lower, registers. Each note should be even in tone and volume, and clearly articulated. If one note sounds different from the others, it’s an indication that the loudspeaker may have a problem at that frequency.

The bottom end should be tight, clean, and “quick.” When it comes to bass, quality is more important than quantity. Poor-quality bass is a constant reminder that the music is being artificially reproduced, making it that much harder to hear only the music and not the loudspeakers. The paradigm of what bass should *not* sound like is a “boom truck.” Those car stereos are designed for maximum output at a single frequency, not articulate and tuneful bass. Unfortunately, more bass is generally an indicator of worse bass performance in low- to moderately priced speakers. A lean, tight, articulate bass is preferable in the long run to the plodding boominess that characterizes inferior speakers.

Listen to kick drum and bass guitar working together. You should hear the bass drum’s dynamic envelope through the bass guitar. The drum should lock in rhythmically rather than seem to lag slightly behind the bass guitar. A loudspeaker that gets this wrong dilutes rhythmic power, making the rhythm sound sluggish, even slower. But when you listen to a speaker that gets this right, you’ll find your foot tapping and hear a more “upbeat” and involving quality in the music.

Midrange coloration is a particularly annoying problem with some loudspeakers. Fortunately, coloration levels are vastly lower in today’s speakers than they were even since the previous edition of this book (2004). Still, there are lots of colored loudspeakers out there. These can be identified by their “cupped hands” coloration on voices, a nasal quality, or an emphasis of certain vowel

sounds. A problem a little higher in frequency is manifested as a “clangy” piano sound. A good loudspeaker will present voices as pure, open, and seeming to exist independently of the loudspeakers. Midrange problems will also make the music sound as if it is coming out of boxes rather than existing in space.

Poor treble performance is characterized by a grainy or dirty sound to violins, cymbals, and vocal sibilants (s and *sh* sounds). Cymbals should not splash across the sound-stage, sounding like bursts of undifferentiated white noise. Instead, the treble should be integrated with the rest of the music and not call attention to itself. The treble shouldn’t sound hard and metallic; instead, cymbals should have some delicacy, texture, and pitch. If you find that a pair of speakers is making you aware of the treble as a separate component of the music, keep looking.

Another thing to listen for in speakers is their ability to play loudly without congestion. The sounds of some loudspeakers will be fine at low levels, but will congeal and produce a giant roar when pushed to high volumes. Listen to orchestral music with loud climaxes—the sound should not collapse and coarsen during loud, complex passages.

Finally, the loudspeakers should “disappear” into the soundstage. A good pair of speakers will unfold the music in space before you, giving no clue that the sound is coming from two boxes placed at opposite sides of the room. Singers should be heard as pinpoint, palpable images directly between the loudspeakers (if that’s how they’ve been recorded). The sonic image of an instrument should not “pull” toward one side or another when the instrument moves between registers. The music should sound open and transparent, not thick, murky, or opaque. Overall, the less you’re aware of the speakers themselves, the better.

Some speakers with less-than-high-end aspirations have colorations intentionally designed into them. The bass is made to be big and fat, the treble excessively bright to give the illusion of “clarity.” Such speakers are usually extremely sensitive, so that they’ll play loudly in comparisons made without level matching. These loudspeakers may impress the unwary in a two-minute demonstration, but will become extremely annoying not long after you’ve brought them home. You’re unlikely to find such products in a true high-end audio store.

Finally, the surest sign that a loudspeaker will provide long-term musical satisfaction at home is if, during the audition, you find yourself greatly enjoying the music and not thinking about speakers at all.

The flaws described here are only the most obvious loudspeaker problems; a full description of what to listen for in reproduced music in general is found in [Chapter 3](#).

Loudspeaker Types and How They Work

Many mechanisms for making air move in response to an electrical signal have been tried over the years. Three methods of creating sound work well enough—and are practical enough—to be used in commercially available products. These are the *dynamic driver*, the *ribbon* transducer, and the *electrostatic* panel. A loudspeaker using dynamic drivers is often called a box loudspeaker because the drivers are mounted in a box-like enclosure or cabinet. Ribbon and electrostatic loudspeakers are called *planar* loudspeakers because they're usually mounted in flat, open panels. The term *transducer* describes any device that converts energy from one form to another. A loudspeaker is a transducer because it converts electrical energy into sound.

The Dynamic Driver

The most popular loudspeaker technology is, without question, the dynamic driver. Loudspeakers using dynamic drivers are identifiable by their familiar cones and domes. The dynamic driver's popularity is due to its many advantages: wide dynamic range, high power handling, high sensitivity, relatively simple design, and ruggedness. Dynamic drivers are also called point-source transducers because the sound is produced from a point in space.

Dynamic loudspeaker systems use a combination of different-size dynamic drivers. The low frequencies are reproduced by a cone *woofer*. High frequencies are generated by a *tweeter*, usually employing a small metal or fabric dome. Some speakers have a third dynamic driver, the *midrange*, to reproduce frequencies in the middle of the audioband.

Despite the very different designs of these drivers, they all operate on the same principle (Fig.6-1). First, the simplified explanation: Electrical current from the power amplifier flows through the driver's *voice coil*. This current flow sets up a magnetic field around the voice coil that expands and collapses at the same frequency as the audio signal. The voice coil is suspended in a permanent magnetic field generated by magnets in the driver. This permanent magnetic field interacts with the magnetic field generated by current flow through the voice coil, alternately pushing and pulling the voice coil back and forth. Because the voice coil is attached to the driver's cone, this magnetic interaction pulls the cone back and forth, producing sound.

On a more technical level, the voice coil is a length of wire wound around a thin cylinder called the voice-coil former. The former is attached to the *diaphragm* (a cone or dome). Electrical current from the amplifier flows through the voice coil, which is mounted in a permanent magnetic field whose magnetic lines of flux cross the gap between two permanent magnets. According to the "right-hand rule" of physics, the circular flow of current through the voice-coil windings generates magnetic forces that are directed along the voice coil's axis. The interaction between the fluctuating field of the voice coil and the fixed magnetic field in the gap produces axial forces that move the voice coil back and forth, carrying the diaphragm with it. The faster the audio signal alternates, the faster the diaphragm moves, and the higher the frequency of sound produced. Dynamic drivers are also called moving-coil drivers, for obvious reasons.

CS 1.5 Woofer





Fig.6-1 Dynamic driver cutaway and cross section (Courtesy Thiel Audio)

Other elements of the dynamic driver include the spider which suspends the voice coil in place as it moves back and forth. The basket is a cast- or stamped-metal structure holding the entire assembly together. (Cast baskets are generally found in higher-quality speakers, stamped baskets in budget models.) A ring of compliant rubber material, called the surround, attaches the cone to the basket rim. The surround allows the cone to move back and forth while still attached to the basket. The maximum distance the cone moves back and forth is called its *excursion*.

Common cone materials include paper, paper impregnated with a stiffening agent, a plastic such as polypropylene, titanium, coated aluminum, or exotic materials such as carbon-fiber, Kevlar, ceramic, beryllium, vapor-deposited diamond, and proprietary composites. Designers use these materials, and sometimes a sandwich of different materials, to prevent a form of distortion called *breakup*. Breakup occurs when the cone material flexes instead of moving as a perfect piston. Because the cone is driven at a small area at the inside (an area the size of the voice coil), the cone tends to flex, which produces nonlinear distortion. Stiff cone materials help prevent breakup. Although all dynamic drivers exhibit breakup at a certain frequency, the competent loudspeaker designer makes sure that a driver is never driven by frequencies that would cause breakup. For example, if a woofer has its first breakup mode at 4kHz, the designer would probably operate the driver only up to about 2kHz, well away from its breakup frequency.

The cone should be lightweight as well as stiff. A lighter cone has less inertia, allowing it to respond faster to transient signals and stop faster after the drive signal has ceased. Think of a bass-drum whack. A large, heavy cone may not be able to move fast enough to reproduce the sound's attack, diminishing the drum's dynamic impact. Similarly, once the drum whack is over, the heavy cone's mass will want to continue moving. Speaker designers therefore search for cone materials that combine high stiffness with low mass. Many of the advances made in loudspeaker design in the past 15 years have been the result of materials research that has yielded lighter yet stiffer diaphragms.

Tweeters work on the same principle, but typically use a 1" dome instead of a cone. Common dome materials include plastic, woven fiber coated with a rubbery material, titanium, aluminum and aluminum alloys, and gold-plated aluminum. A recent trend has been toward making tweeter diaphragms from materials such as beryllium and even diamond. These materials, once found only in a company's

expensive flagship models, have found their way into less costly products. Unlike cone drivers, which are driven at the cone's apex, dome diaphragms are driven at the dome's outer perimeter. Some dome tweeters use Ferrofluid, a liquid cooling agent, to remove heat from the tweeter's small voice coil. The first breakup mode of well-designed, modern dome tweeters is above 25kHz, well beyond the audioband.

Midrange drivers are smaller versions of the cone woofer. Some, however, use dome diaphragms instead of cones.

The Coaxial (Concentric) Driver

A number of loudspeaker manufacturers have created drivers in which a tweeter is positioned inside a midrange driver's cone. Such drivers are called *coaxial* or *concentric* drivers. Notable examples are the Uni-Q from KEF (Fig.6-2), the beryllium midrange and tweeter coincident model made by TAD, and Thiel's corrugated-midrange coaxial driver.



Fig.6-2 KEF's Uni-Q driver positions a dome tweeter inside a midrange cone. (Courtesy KEF)

The concentric driver is a true point source, with none of the nodal cancellation that results when the midrange and tweeter are physically separated on the baffle. The midrange and tweeter output combine perfectly, regardless of the listening distance, listening height, or listening axis. In addition, the midrange cone acts as a *waveguide* for the tweeter, controlling the tweeter's dispersion so that, at the lower end of the tweeter's frequency range, the tweeter's dispersion more closely matches the midrange driver's dispersion. These qualities allow the crossover to be simpler and less intrusive.

Dynamic Compression

When a dynamic driver plays loudly, it must dissipate a tremendous amount of heat from its voice coil. As the voice coil heats up, its electrical resistance increases, reducing the amount of current flow through the voice coil. Because it is current flow through the voice coil that causes the cone to move back and forth, increased voice-coil resistance reduces the amount of sound produced by the driver. In other words, you can keep turning up the volume, but the speaker reaches a point where it won't play any louder.

This phenomenon, called *dynamic compression*, obviously alters musical dynamics. Specifically, loud musical peaks aren't reproduced quite as loudly as they should be. An electroacoustic mechanism is changing the score's dynamic markings; *fff* may be reproduced as *ff*, for example.

Although this distortion of the music's dynamics is a departure from accuracy, dynamic compression introduces a worse form of distortion. The multiple drivers in a three-way dynamic loudspeaker will have different levels at which they begin to compress. If the tweeter exhibits dynamic compression at a lower volume than the midrange and woofer, the sound will be slightly duller on musical peaks. Conversely, if the woofer and midrange compress at a lower level than the tweeter, the sound will grow brighter as it gets louder. In essence, the speaker's tonal balance changes as a function of the music's dynamics. You may not hear this phenomenon directly as a change in tonal balance, but this form of distortion creates an impression of strain on musical peaks that may momentarily draw our attention away from the music. It's just one more cue to the brain that we're not hearing live music. Loudspeaker designers combat dynamic compression by creating drivers with oversize voice coils, physical structures that naturally cool the voice coil, and through innovative voice-coil wire shape and winding techniques. Higher-sensitivity drivers are less prone to dynamic compression, all other factors being equal.

There's a related mechanism by which drivers can distort music. When a driver is producing a loud sound, the ends of its voice coil can momentarily leave the magnetic field in which the coil sits. This robs the driver of its ability to move back and forth faithfully in response to the drive signal, again changing the music's dynamics and creating distortion. This problem is addressed by designing drivers with a very short voice coil and a very long magnetic gap. Although a short voice coil is more susceptible to dynamic compression, it doesn't introduce

this form of nonlinear distortion. Loudspeaker designers must balance these trade-offs.

Finally, a dynamic driver can “bottom out,” heard as a popping or cracking sound from the woofer. The popping sound is caused by the voice-coil former (the bobbin around which the voice coil is wound) hitting the back of the magnet structure. If you ever hear this sound, immediately turn down the volume or the speaker will be damaged.

Problems with Dynamic Drivers

Of all the components in your system, the loudspeaker is the most likely to develop trouble. The most common problem is a blown driver, usually the tweeter—it simply stops working. Tweeters are often destroyed by too much current flowing through their voice coils. The tweeter can't dissipate the resultant heat quickly enough, and its voice coil is burned open.

Another common cause of speaker failure is the buzzing sound produced by a loose voice coil. Too much current through the voice coil melts the glue holding the coil to the former (the bobbin around which wire is wound to make the coil). This loosens turns of wire, which then rub against the magnet to cause the buzzing sound.

Driver mounting bolts can loosen over time and degrade a loudspeaker's performance by allowing the entire driver to vibrate rather than just its diaphragm. Gently tightening these bolts from time to time—particularly when a loudspeaker is new—can improve its sound. Be careful not to overtighten and strip the bolts.

The Electromagnetic Dynamic Driver

The dynamic driver I've described eventually runs up against the laws of physics. Specifically, the magnetic field strength generated by the fixed magnets is limited, which in turn places restrictions on the cone weight, how low in frequency the driver will play, and how sensitive the driver is. A heavy cone goes lower in frequency (all other factors being equal), but requires greater magnetic-field strength surrounding the voice coil to drive it. These limitations are particularly acute in large woofers.

A solution to this physics problem is the *electromagnetic dynamic driver* in which the driver's fixed magnets are replaced with a large coil that functions as an electromagnet. The coil is driven with direct current from an outboard power supply that plugs into an AC outlet. This current flow through the coil creates the magnetic field against which the voice-coil-generated magnetic field pushes and pulls. The electromagnet produces a magnetic field strength in the gap (the area in which the voice coil sits) nearly double that of a conventionally driven woofer. Consequently, the electromagnetic woofer can be heavier (giving it a lower resonant frequency) yet simultaneously more efficient. Moreover, the woofer's bass output can be adjusted by varying the current through the electromagnetic coil. This is accomplished via a control on the outboard supply that drives current through the electromagnetic coil. One can thus adjust the electromagnetic woofer's bass output to better integrate the system into a variety of listening rooms.

An example of the electromagnetic driver is the woofer in the Focal Grande Utopia EM loudspeaker. The system's woofer has a very high sensitivity (97dB for 1W) but very low resonance (24Hz). In other words, the woofer delivers lots of very low bass with very little input power. The price of this performance is the need for the outboard supply that has to be plugged into an AC outlet, along with the sheer weight of the woofer. The EM's 16" woofer weighs 63 pounds, 48 of which is the electromagnetic coil.

The Planar-Magnetic Transducer

The next popular driver technology is the *planar magnetic* transducer, also known as a *ribbon driver*. Although the term *ribbon* and *planar magnetic* are often used interchangeably, a true ribbon driver is actually a subclass of the planar-magnetic driver. Let's look at a true ribbon first.

Instead of using a cone attached to a voice coil suspended in a magnetic field, a ribbon driver uses a strip of material (usually aluminum) as a diaphragm suspended between the north and south poles of two magnets (see [Fig.6-3](#)). The ribbon is often pleated for additional strength. The audio signal travels through the electrically conductive ribbon, creating a magnetic field around the ribbon that interacts with the permanent magnetic field. This causes the ribbon to move back and forth, creating sound. In effect, the ribbon functions as both the voice coil and the diaphragm. The ribbon can be thought of as a stretched-out voice coil.

In all other planar-magnetic transducers, a flat or slightly curved diaphragm is driven by an electromagnetic conductor. This conductor, which is bonded to the back of the diaphragm, is analogous to a dynamic driver's voice coil, here stretched out in straight-line segments. In most designs, the diaphragm is a sheet of plastic with the electrical conductors bonded to its surface. The flat metal conductor provides the driving force, but it occupies only a portion of the diaphragm area. Such drivers are also called quasi-ribbon transducers. [Fig.6-4](#) shows the difference between a true ribbon and quasi-ribbon drivers.




Fig.6-3 A true ribbon driver (After Martin Colloms, *High Performance Loudspeakers*)




Fig.6-4 Cross-section comparison of true ribbon (left) and quasi-ribbon (right) drivers (Courtesy Magnepan)

A planar driver is a true ribbon only if the diaphragm is conductive and the audio signal flows directly through the diaphragm, rather than through conductors bonded to a diaphragm, as in a quasi-ribbon driver. (Despite this semantic distinction, I'll use the term "ribbon" throughout the rest of this section, with the understanding that it covers both true ribbons and quasi-ribbon drivers.)

Ribbon drivers like the one in Figs.6-3 and 6-4 are called *line-source transducers* because they produce sound over a line rather than from a point, as does a dynamic loudspeaker. Moreover, a ribbon's radiation pattern changes dramatically with frequency. At low frequencies, when the ribbon's length is short compared to the wavelength of sound, the ribbon will act as a point source and produce sound in a sphere around the ribbon—just like a point-source woofer. As the frequency increases and the wavelength of sound approaches the ribbon's dimensions, the radiation pattern narrows until it looks more like a cylinder

around the ribbon than a sphere. At very high frequencies, the ribbon radiates horizontally but not vertically. This can be an advantage in the listening room: the listener hears more direct sound from the loudspeaker and less reflection from the sidewalls and ceiling. Reduced wall reflections aid in soundstaging: Pinpoint imaging and the ability to project a concert hall's acoustic signature are hallmarks of good ribbon loudspeakers.

The main technical advantage of a ribbon over a moving-coil driver is the ribbon's vastly lower mass. Instead of using a heavy cone, voice coil, and voice-coil former to move air, the only thing moving in a ribbon is a very thin strip of aluminum. A ribbon tweeter may have one quarter the mass and ten times the radiating area of a dome tweeter's diaphragm. Low mass is a high design goal: the diaphragm can respond more quickly to transient signals. In addition, a low-mass diaphragm will stop moving immediately after the input signal has ceased. The ribbon starts and stops faster than a dynamic driver can, allowing it to more faithfully reproduce transient musical information.

The ribbon driver is usually mounted in a flat, open-air panel that radiates sound to the rear as well as to the front. A loudspeaker that radiates sound to the front and rear is called a *dipole*, which means "two poles." Fig.6-5 shows the radiation patterns of a point-source loudspeaker (left) and a dipolar loudspeaker.




Fig.6-5 A point-source loudspeaker (left) radiates sound in one direction; a dipole loudspeaker (right) radiates sound equally to the front and rear.

Another great advantage enjoyed by ribbons is the lack of a box or cabinet. As we'll see in the section of this chapter on loudspeaker enclosures, the enclosure can greatly degrade a loudspeaker's performance. Not having to compensate for an enclosure makes it easier for a ribbon loudspeaker to achieve stunning clarity and lifelike musical timbres.

A full-range quasi-ribbon loudspeaker is illustrated in [Fig.6-6](#). The large panel extends the system's bass response: when the average panel dimension approaches half the wavelength, front-to-rear cancellation reduces bass output. Consequently, the larger the panel, the deeper the low-frequency extension.

Ribbon loudspeakers are characterized by a remarkable ability to produce extremely clean and quick transients—such as those of plucked acoustic guitar strings or percussion instruments. The sound seems to start and stop suddenly, just as one hears from live instruments. Ribbons sound vivid and immediate without being etched or excessively bright. In addition, the sound has an openness, clarity, and transparency often unmatched by dynamic drivers. (Incidentally, these qualities are shared by ribbon microphones.) Finally, the ribbon's dipolar nature produces a huge sense of space, air, and soundstage depth (provided this spatial information was captured in the recording). Some argue, however, that this sense of depth is artificially produced by ribbon speakers, rather than being a reproduction of the actual recording.

Despite their often stunning sound quality, ribbon drivers have several disadvantages. The first is low sensitivity; it takes lots of amplifier power to drive them. Second, ribbons inherently have a very low impedance, often a fraction of an ohm. Most ribbon drivers therefore have an impedance-matching transformer in the crossover to present a higher impedance to the power amplifier. Good design of the transformer is therefore crucial to prevent it from degrading sound quality.



Fig.6-6 A full-range dipolar ribbon loudspeaker (Courtesy Magnepan)

From a practical standpoint, ribbon-based loudspeakers are more difficult to position in a room. Small variations in placement can greatly change the sound, due primarily to their dipolar radiating pattern. This dipolar pattern requires that ribbon loudspeakers be placed well away from the front wall, and that the front wall be acoustically treated.

Low-profile, ribbon-based loudspeakers with the ribbon top at the same height as the listener's ears will have a radically different treble balance if the listener moves up or down a few inches. That's because ribbon loudspeakers have very narrow vertical dispersion, meaning they radiate very little sound above and below the ribbon at high frequencies. If you sit too high or listen while standing, you'll hear less treble. Some ribbon loudspeakers have a tilt adjustment that allows you to set the correct treble balance for a particular listening height.

Ribbons also have a resonant frequency that, if excited, produces the horrible sound of crinkling aluminum foil. Consequently, the ribbon must be used within strict frequency-band limits. In addition, ribbon drivers are tensioned at the factory for optimum performance. If under too much tension, a ribbon will produce less sound; if under too little tension, a ribbon can produce distortion that sounds like the music is “breaking up.” This is most noticeable on piano; the transient leading edges sound “shattered” and distorted. A sudden increase in ambient temperature can cause a ribbon driver to lose some of its tension and introduce the distortion described. If you hear this sound from your ribbon speakers, contact the manufacturer for advice. The solution may be as simple as turning a few tensioning bolts half a turn.

Ribbon drivers don't necessarily have to be long and thin. Variations on ribbon technology have produced drivers having many of the desirable characteristics of ribbons but few of the disadvantages.

Finally, some loudspeakers use a combination of dynamic and ribbon transducers to take advantage of both technologies. These so-called hybrid loudspeakers typically use a dynamic woofer in an enclosure to reproduce bass, and a ribbon midrange/tweeter (Fig.6-7).

The hybrid technique brings the advantages of ribbon drivers to a lower price level (ribbon woofers are big and expensive), and exploits the advantages of each technology while avoiding the drawbacks. The challenge in such a hybrid system is to achieve a seamless transition between the dynamic woofer and the ribbon tweeter, with no audible discontinuity between the drivers.



Fig.6-7 A dynamic-ribbon hybrid loudspeaker employing a dynamic woofer and a ribbon tweeter (Courtesy Raidho Acoustics)

The Heil Air-Motion Transformer

Designed in the early 1970s by Dr. Oskar Heil, the Heil Air-Motion Transformer (AMT) operates in a completely different way from dynamic or planar drivers. Rather than causing a diaphragm to move back and forth like a piston, the AMT's pleated polyethylene membrane alternately squeezes and expands the air in front of it in response to the audio signal. A conductor is bonded to the diaphragm, and then suspended in a magnetic field. Although ribbons are often pleated for strength, the AMT's pleats are much deeper and are responsible for the squeezing effect. The name Air-Motion Transformer comes from the fact that air is squeezed out of the pleats at five times the speed of the diaphragm's motion. This squeezing effect, and the accompanying acceleration of the air, is analogous to the rapid expulsion of a seed from a piece of fruit that's gently squeezed. The seed's motion is many times faster than the force applied.

The AMT has excellent transient response due to the diaphragm's low mass and very small range of motion, as well as high acoustic output. A 1"-wide AMT produces the same acoustic output as an 8"-diameter driver. A number of companies make a version of the Heil Air-Motion Transformer now that the patents have expired.

The Electrostatic Driver

Like the ribbon transducer, an electrostatic driver uses a thin membrane to make air move. But that's where the similarities end. While both dynamic and ribbon loudspeakers are electromagnetic transducers—they operate by electrically induced magnetic interaction—the electrostatic loudspeaker operates on the completely different principle of electrostatic interaction.

No discussion of electrostatic loudspeakers would be complete without mentioning the classic electrostatic loudspeaker, the Quad ESL-57, created in 1957 by Peter Walker. The ESL-57 revolutionized the standard for transparency on its introduction, and still holds its own more than 50 years later. Many listeners' first exposure to high-quality audio was through an ESL-57, and a large number of contemporary loudspeaker designers still have a pair of ESL-57s on hand as a reference. The ESL-57 doesn't have much low bass, won't play very loudly, and produces a very narrow sweet spot, but when operated within its limitations, it's magical.

In an electrostatic loudspeaker (sometimes called an ESL), a thin, movable membrane—usually made of transparent Mylar—is stretched between two static elements called stators ([Fig.6-8](#)). The membrane is charged to a very high voltage with respect to the stators. The audio signal is applied to the stators, which create electrostatic fields around them that vary in response to the audio signal. The varying electrostatic fields generated around the stators interact with the membrane's fixed electrostatic field, pushing and pulling the membrane to produce sound. One stator pulls the membrane, the other pushes it. This illustration also shows a dynamic woofer as part of a hybrid dynamic/electrostatic system. (A full-range electrostatic loudspeaker is illustrated in [Fig.6-9](#).)




Fig.6-8 An electrostatic loudspeaker employs two stators separated by a membrane that is pulled back and forth by the stators. (Courtesy MartinLogan)



Fig.6-9 A full-range electrostatic loudspeaker (Courtesy MartinLogan)

The voltages involved in an electrostatic loudspeaker are very high. The polarizing voltage applied to the diaphragm may be as high as 10,000 volts (10kV). In addition, the audio signal is stepped up from several tens of volts to several thousand volts by a step-up transformer inside the electrostatic loudspeaker. These high voltages are necessary to produce the electrostatic fields around the diaphragm and stators.

To prevent arcing—the electrical charge jumping between elements—the stators are often coated with an insulating material. Still, if an electrostatic loudspeaker is overdriven, the electrostatic field strips free electrons from the oxygen in the air, making it ionized; this provides a conductive path for the electrical charge. Large diaphragm excursions that are produced to create a loud playing level put the diaphragm closer to the stators and also encourage arcing.

Arcing can destroy electrostatic panels by punching small holes in the membrane. It is a greater problem in humid than in dry climates because moisture makes the air between the stators more electrically conductive.

Electrostatic panels are often divided into several smaller panels to reduce the effects of diaphragm resonances. Some panels are curved to reduce the lobing effect (uneven radiation pattern) at high frequencies. Lobing occurs when the wavelength of sound is small compared to the diaphragm. Lobing is responsible for electrostatics' uneven high-frequency dispersion pattern, which *Stereophile* magazine founder J. Gordon Holt has dubbed the *vertical venetian-blind effect*, in which the tonal balance changes rapidly and repeatedly as you move your head from side to side.

Electrostatic panels are of even lighter weight than planar-magnetic transducers. Unlike the ribbon driver, in which the diaphragm carries the audio-signal current, the electrostatic diaphragm need not carry the audio signal. The diaphragm can therefore be very thin, often less than 0.001". Such a low mass allows the diaphragm to start and stop very quickly, improving transient response. And because an electrostatic panel is driven uniformly over its entire area, it is less prone to breakup. Both the electromagnetic planar loudspeaker (a ribbon) and the electrostatic planar loudspeaker enjoy the benefits of limited dispersion, which means less reflected sound arriving at the listening position. Like ribbon speakers, electrostatic speakers also have no enclosure to degrade the sound. Electrostatic loudspeakers also inherently have a dipolar radiation pattern. Because the diaphragm is mounted in an open panel, the electrostatic driver produces as much sound to the rear as to the front. Finally, the electrostatic speaker's huge surface area confers an advantage in reproducing the correct sizes of instrumental images.

In the debit column, electrostatic loudspeakers must be plugged into an AC outlet to generate the polarizing voltage. Because the electrostatic is naturally a dipolar radiator, placement in the room is more crucial to achieving good sound. The electrostatic speaker needs to be placed well out into the room and away from the front wall to realize a fully developed soundstage. Electrostatics also tend to be insensitive, requiring large power amplifiers. The load impedance they present to the amplifier is also more reactive than that of dynamic speakers, further taxing the power amp. (Reactance is described later in this chapter.) Nor will they play as loudly as dynamic speakers; electrostatics aren't noted for their dynamic impact, power, or deep bass. Instead, electrostatics excel in transparency, delicacy, transient response, resolution of detail, stunning imaging, and overall musical coherence.

Electrostatic loudspeakers can be augmented with separate dynamic woofers or a subwoofer to extend the low-frequency response and provide some dynamic impact. Other electrostatics achieve the same result in a more convenient package: dynamic woofers in enclosures mated to the electrostatic panels. Some of these designs achieve the best qualities of both the dynamic driver and electrostatic panel. (An example of a hybrid electrostatic/dynamic loudspeaker is illustrated in Fig.6-10.) Audition such hybrid speakers carefully; they sometimes exhibit an audible discontinuity at the transition frequency at which the woofer leaves off and the electrostatic panel takes over. Listen, for example, for a change in a piano's timbre, bloom, projection, and image size as it is played in different registers. Acoustic bass in jazz is also a good test of woofer/panel discontinuity in dynamic/electrostatic hybrid loudspeakers.

One great benefit of full-range ribbons and full-range electrostatics is the absence of a crossover; the diaphragm is driven by the entire audio signal. This prevents any discontinuities in the sound as different frequencies are reproduced by different drivers. In addition, removing the resistors, capacitors, and inductors found in crossovers greatly increases the full-range planar's transparency and harmonic accuracy. Even hybrid planars put the crossover frequency between the dynamic woofer and the planar panel very low (below 800Hz), so there's no discontinuity between drivers through most of the audioband.

Finally, the large diaphragms of electrostatic and ribbon drivers are gently driven over their entire surface areas, rather than forcefully over the relatively small voice-coil area of a dynamic driver. This high force over a small area contributes to the dynamic driver's breakup described earlier, a phenomenon less likely to occur with a large planar diaphragm.



Fig.6-10 A hybrid electrostatic loudspeaker employs a dynamic woofer mated to an electrostatic midrange/tweeter panel. (Courtesy MartinLogan)

The Dipolar Radiation Patterns of Ribbons and Electrostatics

Because planar loudspeakers (ribbons and electrostatics) are mounted in an open frame rather than an enclosed box, they radiate sound equally from the front and back. As we saw earlier, the term *dipolar* describes this radiation pattern. This is contrasted with point-source loudspeakers, whose drivers are mounted on the front of a box. Point-source loudspeakers are usually associated with dynamic drivers, but any type of driver in an enclosed cabinet qualifies as a point-source speaker.

The dipolar radiation patterns of ribbons and electrostatics make using them very different from point-source loudspeakers. Because they launch just as much energy to the rear as they do to the front, positioning a dipole is more crucial, particularly their distance from the front wall. Dipoles need a significant space behind them to work well. In addition, the front wall's acoustic properties have a much greater influence over the sound. A few attempts have been made to absorb a ribbon's rear energy inside an enclosure to make placement easier, but these have been largely unsuccessful.

Ribbons and electrostatics also tend to have much narrower radiation patterns than point-source loudspeakers. That is, they disperse less sound to their sides. Consequently, the listener hears much more of the speaker's direct sound and much less sound reflected from the floor, ceiling, and sidewalls. This is particularly true in the midrange and treble, where the radiation pattern becomes cylindrical. The result is often a greatly improved transparency and more natural timbre. (The effects of these reflections are discussed in [Chapter 14](#).) The negative effects of this tendency of planar loudspeakers to “beam” at high frequencies is a disturbing shift in the tonal balance with small motions of the head from side to side—the “vertical venetian-blind effect” described earlier. This tendency can be minimized with a curved diaphragm (seen in [Figs.6-9](#) and [6-10](#)) or by electrically segmenting the diaphragm.

Bipolar and Omnidirectional Loudspeakers

We've seen that a diaphragm mounted in an open panel will naturally radiate sound equally to the front and rear. These front and rear waves are *out of phase* with each other. That is, as the diaphragm moves forward, it creates a positive pressure wave in front of the diaphragm, and a negative pressure wave (called a *rarefaction*) behind the diaphragm.

Some loudspeaker designers see the advantages of front- and rear-launched sound, but believe that the front and rear waves should be in phase with each other. Because this is impossible to achieve with a diaphragm mounted in an open panel, these designers employ a second set of dynamic drivers on the rear of the speaker enclosure. This design, called a *bipolar* loudspeaker, has a radiation pattern similar to that of a dipole, but the front and rear waves are now in phase. Other designers add just a rear-firing tweeter, which, they contend, produces a smoother treble response at the listening position and better reproduces a sense of spaciousness and ambience. To recap: In a dipolar loudspeaker, the front and rear waves are out of phase with each other. In a bipolar loudspeaker, the front and rear waves are in phase.

A bipolar loudspeaker produces a radiation pattern slightly different from that of a dipole. The dipole's out-of-phase front and rear waves cancel at the sides of the speaker, narrowing their radiation pattern. A bipolar speaker's in-phase waves don't exhibit this cancellation, and thus have wider dispersion.

Advocates of dipoles and bipoles believe that reflecting sound from the wall behind the speakers increases the sense of musical realism because we hear such reflected sound in live music. In addition, point-source speakers (those firing in one direction) exhibit frequency-response changes off-axis (to the speaker's sides). The sound reflected from the sidewalls has a different spectral balance compared with the direct sound from the loudspeaker. This confuses the ear/brain, and provides yet another clue that the sound isn't real. Because dipole and bipole loudspeakers distribute a more uniform spectral balance around the room, aficionados of such designs believe that they produce a more musically natural sound.

An extension of this idea is the omnidirectional loudspeaker which produces a spherical dispersion pattern; that is, it radiates sound in all directions. Proponents claim that omnidirectional speakers produce a ratio of direct to reflected sound that more closely matches that of live music. Omnidirectional loudspeakers produce a large and expansive soundstage, but not "pinpoint"

imaging. Fig.6-11 shows an omnidirectional speaker, the MBL X-Treme. Note also that this design employs a novel driver technology called a Radialstrahler, in which strips of thin metal are arranged to form an ovoid. The strips are connected at the top and bottom to magnetic motors. These motors, driven by the audio signal, cause the strips (called “petals”) to flex, producing sound.



Fig.6-11 An omnidirectional loudspeaker radiates sound in all directions. (Courtesy MBL)

The Horn Driver

One of the first loudspeaker technologies, the *horn*, has enjoyed a resurgence in popularity in the last 20 years. The horn loudspeaker employs a small dynamic driver mounted in the throat (narrow end) of a horn structure, which more efficiently couples the driver's diaphragm to the air (Fig.6-12).

Any type of driver can be horn loaded, but it is uncommon to find true horn loading in woofers because of the enormous size required. The lower the frequency at which the horn is designed to operate, the larger the horn must be. Consequently, horn loading is generally confined to midrange drivers and tweeters. Full-range horn-based systems commonly use a conventionally loaded woofer to reproduce bass.

Horn-loading a driver confers many performance benefits. First, horns have extraordinarily high sensitivity, and can be driven to very high volume with just a few watts of power. It's not unusual for a horn-loaded system to have a sensitivity of 100dB or more. This brings us to the second benefit of horns: they can be driven by very small power amplifiers, often those producing just 10Wpc or less. Third, the excursion (back-and-forth motion) of the cone or dome diaphragm can be an order of magnitude less than in a direct-radiating driver. This allows the driver to operate in the linear range of its motion at all times, greatly reducing distortion. The electrical and magnetic forces involved in moving a horn-loaded diaphragm back and forth are about one-tenth those of a direct-radiating driver. As a result of these attributes, horn-loaded systems have state-of-the-art dynamics and lifelike re-creation of music's transient nature. The sharp attack of a snare drum, for example, is reproduced with a stunning sense of realism. This quality, known as "jump factor," gives horn-loaded loudspeakers a lifelike sound unmatched by any other loudspeaker technology.

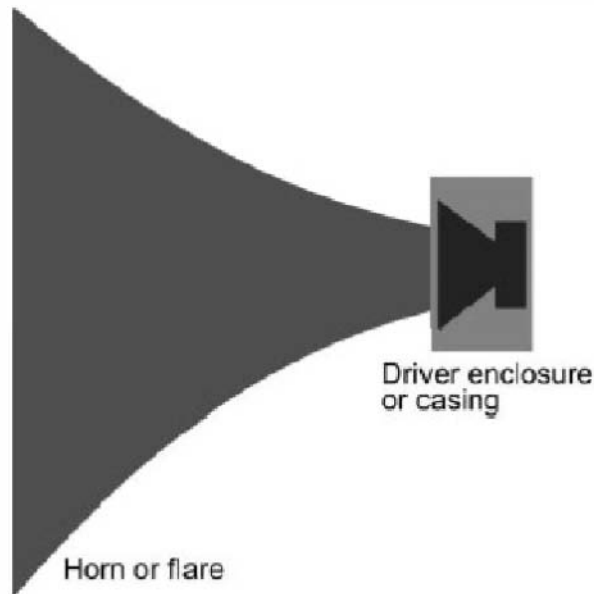


Fig.6-12 Horn-loading a driver more efficiently couples the driver's diaphragm to the air, increasing sensitivity.

Now the bad news. Horn loading often introduces large tonal colorations. Try this experiment: read the previous sentence aloud, and then read it again, this time with your hands cupped around your mouth to make a horn. For many listeners, the horn-loaded system's colorations, particularly through the midrange, are a deal-breaker, outweighing all of the horn's advantages.

There are, however, a few horn-loaded systems that do not sound colored in the slightest. These systems are very large, exotic, and extremely expensive, the result of the extraordinary manufacturing required. The horn-loaded system in [Fig.6-13](#) is an example. In this five-way system, three of the horns are machined out of solid blocks of aircraft-grade aluminum; the top horn, which handles lower-midrange frequencies, is constructed from thick aluminum reinforced with 56 hand-welded ribs. Note the lower-midrange horn's massive size, necessary to maintain the horn-loading effect down to 120Hz, the frequency at which it is crossed over to the conventionally loaded woofer. If a horn's mouth isn't large enough relative to the wavelength being reproduced, the horn's natural acoustic amplification effect diminishes. At this point, the driver operates like a conventionally loaded driver and sensitivity drops off rapidly.

For the technically minded, I've included a more in-depth description of how a horn structure acts as a natural acoustic amplifier. The horn speaker operates on the same principle as a megaphone: The sound source (the driver, in the case

of the horn-loaded speaker) is mounted in the narrow end of a tube that exhibits an exponential flare. This narrow throat restricts air motion and thus presents a high acoustical impedance to the driver diaphragm. Consequently, the diaphragm can develop very high pressures in the throat with very little excursion. As this high-pressure wavefront travels down the horn, the horn's exponentially increasing size gradually converts this high pressure into increasingly lower pressure, until it reaches the horn's mouth. At the mouth, the waveform is of lower pressure but higher amplitude. The horn's shape essentially converts the high-pressure, low-displacement wave-front over a small area (in the throat) to a low-pressure, high-amplitude wavefront to a larger area (at the mouth). The horn has been likened to an acoustic impedance-matching transformer that more efficiently couples the driver diaphragm to the air.



Fig.6-13 A massive horn loudspeaker (Courtesy Magico)

A related technology is the *compression driver*, a dynamic transducer used in horn loudspeakers. The compression driver gets its name from the fact that it uses a diaphragm to “compress” the air through an aperture that is very small relative to the diaphragm’s size. When a compression driver is mounted in the throat of a horn, the combination is capable of extremely high sensitivity—as high as 115dB SPL for 1W of power. The compression driver employs massive magnets relative to the diaphragm size, and consequently, is quite heavy. Compression drivers hand-built by Japanese artisans can cost as much as \$15,000 each. You can see such compression drivers in [Fig.6-13](#); they are the small cylindrical structures at the throats of the three horns.

Waveguides

A tweeter is sometimes mounted in a recessed flare in the baffle, a structure called a waveguide. The waveguide tailors the tweeter's dispersion so that it more closely matches the midrange driver's dispersion at the crossover frequency. This technique avoids a discontinuity at the midrange-tweeter transition that could be audible.

To understand how a waveguide works, you need to know that a midrange driver has a very wide dispersion at the lower end of its frequency range. At higher frequencies (shorter wavelengths relative to the driver size), the midrange driver's dispersion is narrower. Similarly, the tweeter's dispersion is wide at the low end of its frequency range (because the wavelengths are long relative to the driver size) and narrows as the soundwaves it is reproducing increase in frequency. The waveguide physically restricts the tweeter's dispersion so that it more closely matches the midrange driver's dispersion at the crossover point. A waveguide also has a slight horn-loading effect (described earlier in this chapter) that increases the tweeter's sensitivity. That is, a tweeter in a waveguide requires less amplifier power to produce the same sound-pressure level. This has the added benefit of reducing the stress on the tweeter; the diaphragm doesn't have to move as far, and the voice coil doesn't have to handle as much current. These benefits give the loudspeaker designer the option of setting a lower crossover frequency between the midrange driver and the tweeter. The waveguide's shape and depth determine the degree of this horn-loading effect. Taken to an extreme, the waveguide can introduce the problems of horns described earlier in this chapter, which is why waveguides are generally shallow and have a gentle flare, as seen in [Fig.6-14](#).



Fig.6-14 A tweeter mounted in a waveguide (Courtesy Revel)

The Plasma Driver

Have you ever heard the sound of electricity arcing, or the crack of thunder? Both sounds are created by the same phenomenon: An electrical charge jumping between two points through air causes the air around the arc to become ionized (turned into plasma, in physics parlance). The transient expansion and collapse of the plasma create an acoustic wave that we hear as sound. Loudspeaker designers have capitalized on this phenomenon by creating the *plasma driver*, in which an electric arc is modulated by the audio signal. The massless driver is the holy grail of loudspeaker design. It's the moving mass of speaker cones, domes, and panels that introduces a whole host of distortions—distortions that designers devote an inordinate amount of time, money, and effort to combating. The massless driver has no inertia, no resonance, no breakup modes, no nonlinear motion, no temperature-dependent dynamic compression, and no magnetic eddy currents. The list could go on and on.

The catch is that the plasma driver requires a specialized amplification circuit integral to the driver, along with a power supply that produces a very high voltage to create the ionization of the surrounding air. The plasma driver is capable of reproducing only high frequencies, confining its use to tweeters. This fact presents the speaker designer with a dilemma: how to integrate the sound of a massless tweeter with a relatively massive moving-coil midrange driver to create a seamless overall sound. It's a challenge that few have attempted. But when well executed, a loudspeaker with a plasma tweeter sounds stunningly magical. The treble has an ethereal quality, as though the music exists independently of any electromechanical contrivance.

Surprisingly, the plasma driver was invented in 1900; its inventor called it the singing arc. The technology was forgotten until the 1950s, when it was commercialized by DuKane and called the Ionovac. The Ionovac design was bought by the speaker manufacturer Lansche, which produces several modern loudspeaker systems based on it.

The Full-Range Driver

There's one final class of driver we'll look at, the *full-range driver*. Rather than split up the audioband into bass, midrange, and treble frequencies as in a three-way loudspeaker, the full-range speaker attempts to reproduce the entire audioband with a single driver. The advantages are: no crossover parts in the signal path, no interaction between drivers, no phase anomalies introduced by the crossover, and no sonic discontinuities between drivers of different sizes, shapes, and materials. The trade-off is often less-good performance at the frequency extremes (low bass and high treble) and reduced power handling. The latter problem can be addressed with multiple full-range drivers in the same enclosure, all of which reproduce the same signal. Sonically, full-range loudspeakers can have a wonderful coherence and purity in the midrange. This is true even of relatively inexpensive full-range loudspeakers, making them a good choice for a desktop system or in a small room, particularly if you value midrange transparency and resolution over the last measure of bass impact and dynamics.

Loudspeaker Enclosures

The enclosure in which a set of drivers is mounted has an enormous effect on the loudspeaker's quality of reproduced sound. In fact, the enclosure is as important as the drive units themselves. Designers have many choices in enclosure design, all of which affect the reproduced sound.

In addition to the very different enclosure types described in this section, any speaker cabinet will vibrate slightly and change the sound. The ideal enclosure would produce no sound of its own, and not interfere with the sound produced by the drivers. Inevitably, however, some of the energy produced by the drive units puts the enclosure into motion. This enclosure vibration turns the speaker cabinet into a sound source of its own, which distorts the music.

One of the factors that makes current high-end loudspeakers so much better than mass-produced models is the extreme lengths to which some manufacturers go to prevent the enclosure from vibrating. Mass-market manufacturers generally skimp on the enclosure because, to the uninformed consumer, it adds very little to the product's perceived value.

Let's look at why a loudspeaker needs an enclosure, and survey the most popular enclosure types.

An enclosure around a woofer is required to prevent the woofer's rear-radiated wave from combining with its front-radiated wave. When the woofer moves forward, a positive pressure wave is launched at the front, and a rarefaction (negative pressure) wave of equal intensity is launched at the rear. Because low frequencies can refract (bend) around small objects such as a woofer, the rear wave will combine with the front wave and cancel the sound, resulting in reduced bass output. Putting an enclosure around the woofer prevents this cancellation and allows the loudspeaker to better generate low frequencies. Although this phenomenon occurs with planar loudspeakers, their large size lowers the frequency at which this cancellation occurs. That's why full-range planar loudspeakers don't go very low in the bass, and are often augmented by a dynamic cone woofer in a traditional enclosure.

In addition to preventing this front-to-back cancellation, the enclosure should also present the optimal environment for the drivers, particularly the woofer. This is called *loading*. Various loading techniques have evolved to optimize the way the woofer operates in the cabinet. The most common techniques are the *infinite baffle* (also called a *sealed* cabinet, or *air suspension* in smaller enclosures),

and the *reflex* or *ported* enclosure. Less common are the *transmission line*, *Isobarik*, and *finite baffle* loading techniques.

These different loading techniques affect the speaker's bass response. A woofer in an enclosure will go only so low in the bass; as the frequency drops below a certain point, so does the woofer's output. The point at which the woofer's output is down by 3dB is called the low-frequency cutoff point. Below this frequency, the woofer produces less and less sound as the signal driving the woofer decreases in frequency—a phenomenon called *rolloff*. Rolloff is why not every speaker can reproduce 20Hz bass. The loading technique chosen by the speaker designer greatly affects the woofer's low-frequency cutoff point and how steeply this rolloff occurs.

Let's look at the advantages and drawbacks of the most popular loading techniques.

Infinite-Baffle Loading

The *infinite baffle* simply seals the enclosure around the driver rear to prevent the two waves from meeting. In theory, such an enclosure approximates a baffle (the surface on which the drivers are mounted) of “infinite” size; in fact, it wraps around the driver. The air inside the sealed enclosure acts as a spring, compressing when the woofer moves in and creating some resistance to the woofer’s motion. If this resistance to woofer motion by the compressed air is great enough, the infinite baffle can be called *air suspension* (we can dispense with the technical details of where this transition occurs). Although all air-suspension speakers are infinite baffle, not all infinite baffles are air suspension. Nonetheless, all sealed-enclosure speakers have a low-frequency rolloff slope of 12dB per octave; this means that, one octave below the system’s resonant frequency, the output will be reduced by 12dB. This rolloff is relatively gradual, meaning you’ll still hear some lower bass.

Reflex Loading

I mentioned earlier that a woofer in a box produces just as much sound inside the box as outside. The *bass-reflex* enclosure exploits this fact, channeling some of this bass energy out of the cabinet **and into** the listening room. You can instantly spot a reflex-loaded loudspeaker by the hole, or vent, on the front or back of the enclosure. Reflex-loaded loudspeakers are also called *vented* designs.

Reflex loading extends the low-frequency cutoff point by venting some of the energy inside the enclosure to the outside. That is, a reflex-loaded woofer will go lower in the bass than one in a sealed enclosure, other factors being equal. The designer can't, however, just cut a hole in the cabinet. Instead, a duct inside the cabinet delays the woofer's rear-firing wave so that it emerges in phase with the front-firing wave. The rear wave channeled outside the cabinet by the port thus reinforces rather than cancels the front wave. A loading material such as long-hair wool or fiberglass batting is sometimes put in the duct to smooth any peaks in the sound emerging from the port.

Reflex loading has three main advantages. First, it increases a speaker's maximum acoustic output level—it will play louder. Second, it can make a speaker more sensitive—it needs less amplifier power to achieve the same volume. Third, it can lower a loudspeaker's cutoff frequency—the bass goes deeper. Note that these benefits are not available simultaneously; the acoustic gain provided by reflex loading can be used either to increase a speaker's sensitivity or to extend its cutoff frequency, but not both.

Although a reflex system's low-frequency cutoff point is lower than that of a sealed system (assuming the same woofer and enclosure volume), the reflex system's bass rolls off at a much faster rate. Specifically, a reflex system rolls off at 24dB/octave compared to the sealed enclosure's 12dB/octave rolloff. If all other factors are equal, the reflex-loaded system maintains flat bass response down to a lower frequency, but then the bass output drops off more quickly than it does in a sealed system. A comparison of low-frequency cutoff points and rolloff slopes is illustrated in [Fig.6-15](#).

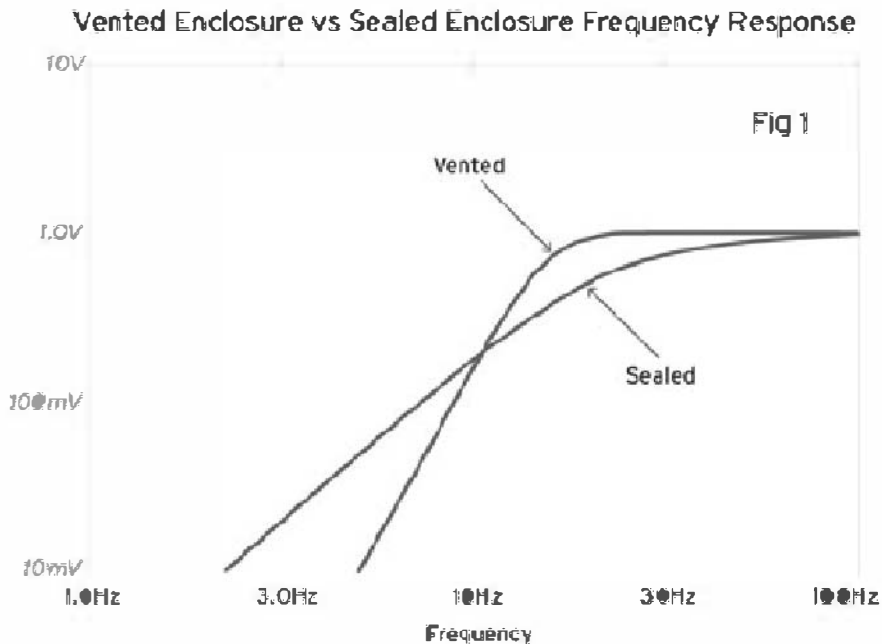


Fig.6-15 Frequency-response characteristics of vented and sealed enclosures

Subjectively, the sealed loudspeaker's higher cutoff frequency and gradual rolloff provide a more satisfying feeling of bass fullness than the reflex system's lower cutoff frequency and steeper rolloff. Some designers have found that a loudspeaker's -10dB point is the most reliable indicator that a loudspeaker's subjective bass fullness and extension because it takes into account not only the low-frequency cutoff point, but also the steepness of the rolloff.

Reflex loading is more often used in small enclosures, which trade low-frequency extension and sensitivity for compactness. It's easier to get more bass from a small reflex enclosure than from a sealed one. However, a small reflex enclosure can still achieve some of the extension and sensitivity of a larger cabinet.

Some controversy exists over the merits of reflex loading. Some designers believe that reflex systems are fundamentally wrong. Reflex systems in general have gotten a bad name, because of both inherent limitations and poor implementation in certain products. Small, inexpensive loudspeakers are often reflex loaded to increase their sensitivity.




Fig.6-16 Transient-response comparison between vented and sealed enclosures

Ineptly implemented reflex systems tend to be boomy, slow, and poorly defined in the bass. In addition, you can sometimes hear the sound emerging from the port, a phenomenon perfectly (and onomatopoeically) described as *chuffing*. Critics also point to the reflex system's poor transient performance compared with that of a sealed loudspeaker. Specifically, a woofer in a vented enclosure will tend to keep moving after the drive signal has stopped, as shown in Fig.6-16. This difference in transient response can be manifested as sluggish sound in, for example, a kick drum. A sealed enclosure has better transient response and better bass definition, but at the cost of lower sensitivity and less deep-bass extension. Nonetheless, many very-high-quality loudspeakers have been based on ported designs.

Passive Radiators

A variation on the reflex system is the *passive radiator*, also called an *auxiliary bass radiator*, or *ABR*. This is usually a flat diaphragm, with no voice coil or magnet structure, that cannot produce sound on its own. Instead, the diaphragm covers what would have been the port in the reflex system, and moves in response to variations in air pressure inside the cabinet caused by the woofer's motion. The passive radiator smoothes out any response peaks, and eliminates a ported system's wind noise (the chuffing mentioned earlier) and port resonances.

Transmission-Line Loading

Another loading technique is the *transmission line*. In a transmission line, the rear wave from the woofer is channeled through a folded labyrinth or duct—the transmission line—filled with absorbent material such as wool. The end of the transmission line appears at the outside of the enclosure, just as in a reflex system. The acoustic energy traveling down the transmission line is absorbed, with the goal of no acoustic output at the end of the transmission line. In theory, the woofer’s rear wave is dissipated in the line. This isn’t the case in practice: most transmission lines have some output at the end of the line, causing them to function partially as reflex systems. Because the line must be at least as long as the lowest wavelength to be absorbed, transmission-line enclosures are usually very large. Proponents of transmission lines argue that the springiness of sealed enclosures is not the ideal environment for the woofer. The transmission line removes the spring action of the sealed enclosure without the disadvantages of reflex designs.

Transmission-line loudspeakers have extremely deep extension and the ability to deliver very loud and very clean bass with an effortless quality not heard from other loading designs. They aren’t as popular as their sound quality would suggest they should be, mostly because of their large size.

British loudspeaker manufacturer Bowers and Wilkins (B&W) has developed a variation on transmission-line loading the company calls Nautilus. Each driver in a Nautilus system is mounted in a tapered or coiled (thus the Nautilus name) enclosure that absorbs the driver’s rear wave. The rear wave cannot interact with the drivers, nor can it form standing waves inside the enclosure. An exploded view of a B&W loudspeaker using Nautilus technology is shown in [Fig.6-17](#). Note also the spherical midrange enclosure, which reduces cabinet diffraction. This “boxless” approach to loudspeaker design is an attempt to reduce the enclosure’s contribution to the reproduced sound.




Fig.6-17 B&W's Nautilus technology absorbs the drivers' rear-firing energy via a computer-optimized tapered-tube enclosure. (Courtesy Bowers & Wilkins Loudspeakers)

Isobarik Loading

A fourth type of loading, used primarily by Linn Products, is called *Isobarik* (Linn's trade name), or by the more descriptive generic term *constant pressure chamber* (see Fig.6-18). In the Isobarik system, a second woofer is mounted directly behind the first woofer and driven in parallel with it. As the front woofer moves back, so does the second; this maintains a constant pressure inside the chamber between the two woofers. This technique offers deeper low-frequency extension, higher power handling, greater linearity, and reduced standing-wave reflections inside the enclosure. Isobarik loading reduces sensitivity because the amplifier must drive two woofers, although only one produces acoustic output. Two woofers mounted in an Isobarik configuration can be modeled as a single woofer whose greater mass and compliance deliver deep bass in an enclosure half the usual size.




Fig.6-18 Isobarik loading (Courtesy Martin Colloms, *High Performance Loudspeakers*)

The Finite Baffle

The only examples of the *finite-baffle*, or open-panel, loading technique are planar loudspeakers such as the full-range ribbon system illustrated in [Fig.6-6](#). The finite baffle has a rolloff of only 6dB/octave until the driver's free-air resonance frequency, then drops at the rate of 18dB/octave.

System Q

Just as a struck bell produces a certain pitch, a woofer in an enclosure will naturally resonate at some frequency. The nature of that resonance is an important characteristic of the loudspeaker that greatly influences its sound. The term Q, for “quality factor,” is a unitless number that expresses how a woofer resonates in an enclosure.

Specifically, a loudspeaker’s Q equals the resonant peak’s center frequency divided by the peak’s bandwidth. A woofer that “rings” (resonates) over a very narrow frequency band is said to have a higher Q than a woofer that resonates less severely over a wider band of frequencies. The narrower the resonance, the higher the Q.

The woofer has its own resonant Q, which is modified by the enclosure’s Q. These resonances combine and interact to reach the system Q, which usually falls between 0.7 and 1.5, as shown in Fig.6-19. A Q of less than 1 is considered *overdamped*, while a Q of more than 1 is *underdamped*. You’ll sometimes hear a loudspeaker described as having subjectively “underdamped bass,” which means the bass is full and warm but lacks tautness. Technically, these terms refer to the system’s anechoic response (the speaker’s response in a reflection-free room), specifically whether the response is up or down at the resonant frequency. A “critically damped” system having a Q of 0.5 provides perfect transient response, with no detectable overhang. That is, the woofer stops moving the instant the drive signal stops. The higher the Q, the more the woofer rings.

Subjectively, an underdamped alignment has lots of bass but lacks tautness, has poor pitch definition, and tends to produce “one-note” bass. An overdamped alignment produces a very tight, clean, but decidedly lean bass response. An overdamped loudspeaker has less bass, but that bass is of higher quality than the bass from an underdamped system. Overdamped speakers tend to satisfy intellectually by resolving more detail in the bass, but often lack the bass weight and power that viscerally involve your whole body in the music. Most loudspeaker designers aim for a Q of about 0.7, to reach a compromise between extended bass response (down only 3dB at resonance) and good transient response (very slight overhang). Some designers maintain that a Q of 0.5 is ideal, and that a higher Q produces bass of poorer quality.

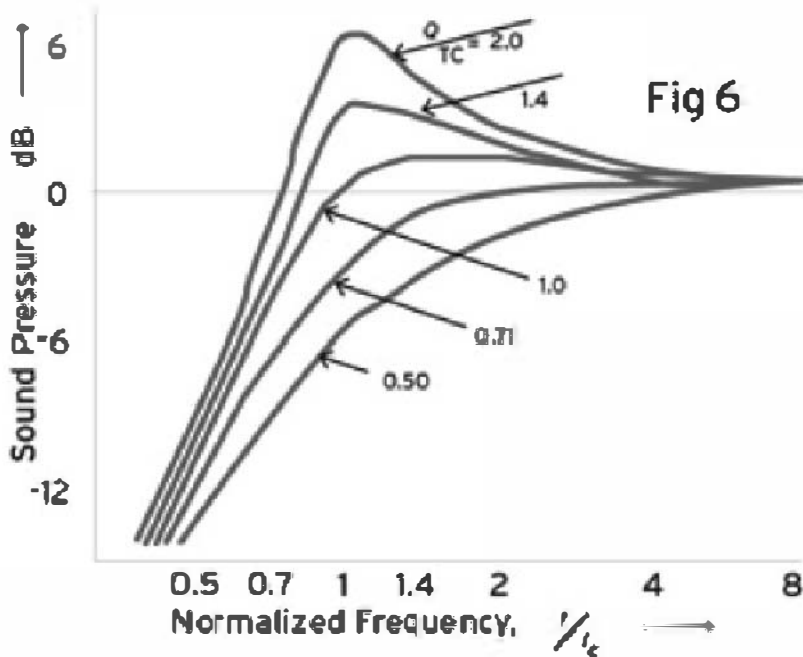


Fig.6-19 A comparison of system Qs (After Martin Colloms, *High Performance Loudspeakers*)

Mass-market loudspeakers are virtually always underdamped (high Q) so that the unwary will be impressed by the loudspeaker's big bottom end. An example of absurdly high Q is the "boom truck" that produces a big bass impact but fails to resolve pitch, dynamic nuances, or any semblance of musical detail. That boom you hear is the woofer resonating in its enclosure over a very narrow frequency range—the antithesis of what we want in a high-end loudspeaker.

Powered and Servo-Driven Woofers

Loudspeaker designers are increasingly choosing to include a power amplifier within the loudspeaker to drive the woofer cone. In such designs, the external power amplifier is relieved of the burden of driving the woofer, and the designer has more control over how the woofer behaves. For example, equalization can be applied to extend the bass response.

In the most notable and successful examples of a powered-woofer system, the Vandersteen Model 5A and Model 7, the bass is rolled off by an external crossover before it reaches the loudspeaker, and a complementary boost restores flat response to the signal driving the woofer. This technique eliminates crossover parts (specifically, the large series inductor) between the internal amplifier and the woofer cone, which makes it easier for the amplifier to more tightly control the woofer's motion. This is audible as better-defined bass reproduction. This loudspeaker also employs an 11-band equalizer to minimize room-induced frequency-response variations below 120Hz. As you can see, a loudspeaker employing an integral power amplifier has some decided advantages.

Some designers don't try to juggle these laws of physics to produce the most musically satisfying compromise in bass performance; instead, they take brute-force control of woofer movement with the *servo-driven woofer*. A servo woofer system consists of a woofer with an accelerometer attached to its voice coil, and a power amplifier dedicated to the woofer. An accelerometer is a device that converts motion into an electrical signal. The accelerometer sends a signal back to the woofer amplifier, telling the woofer amplifier how the woofer cone is moving. The woofer amplifier compares the drive signal to the cone's motion and interprets any difference as distortion. The woofer amplifier can then change the signal driving the woofer so that the woofer cone's motion perfectly matches the drive signal.

For example, the inertia in a woofer would cause it to continue moving after a bass-drum whack. The woofer servo amplifier would see the cone moving and instantly stop that motion. In fact, if you try to gently press in the cone of a servo-driven woofer with the amplifier connected without playing music, you'll find that the cone doesn't move. Because there's no drive signal, the servo amplifier knows the cone shouldn't be moving and locks it in place.

A servo-driven woofer doesn't follow the same rules as a woofer in an enclosure with no servo drive. Instead of rolling off in the bass below the woofer's

resonant frequency, a servo-driven woofer can have flat response down to very low frequencies. The woofer's natural rolloff is counteracted by equalization in the servo amplifier that forces the woofer to keep going lower and lower in frequency. The servo amplifier simply pumps more and more current through the woofer's voice coil as the frequency descends, which increases the woofer's back-and-forth excursion to produce more sound at low frequencies. The lower the frequency, the farther the woofer cone must move back and forth to maintain the same volume.

This puts enormous stress on both the woofer and the servo amplifier. The woofer's voice coil must be able to handle the large amount of current flow, and the woofer cone is susceptible to flexing because it is driven over an area the size of the voice coil—the outside area of the cone may not move exactly in step with the area near the voice coil.

Servo systems therefore must be designed to withstand such rigors. Servo amplifiers are often large and powerful, with huge power transformers and many output transistors. The woofer cone may be made of a trilaminate material that sandwiches a damping compound between two metal diaphragms. This structure ensures that the cone moves as a perfect piston, even when driven hard. Oversized voice coils provide high power handling, and the woofer's structure accommodates high excursion.

A loudspeaker that has servo-driven woofers may use many woofers operating together. An array of woofers reduces the contribution each must make to the sound, reducing the demands placed on individual woofers. Not every woofer in such an array will have an accelerometer; one accelerometer may serve four woofers, the servo amplifier assuming that the other three woofers are behaving identically.

Servo-driven woofers are characterized by a sense of slam and impact in the bottom end. There's a feeling of suddenness to bass transients that sounds closer to live music. Moreover, the bass stops just as quickly, with no overhang, slowness, or blurring of the music's dynamic structure. Kick drum has a crispness and punch rather than sounding like a dull thud. I also find that servo-driven woofers have superb rendering of pitch and precise articulation of each note. Servo-driven woofers can also produce deep bass from a small enclosure.

Servo-driven woofer systems are uncommon because they are more expensive than "cones-in-a-box" loudspeakers, and require specialized expertise to get them to work properly. Keep in mind that any loudspeaker system that uses a dedicated woofer amplifier and a separate amplifier to drive the midrange and tweeter

enjoys all the advantages conferred by bi-amping, such as greater ability to play loudly without strain.

Enclosure Resonances

A casual acquaintance once tried to impress me with how great his brother's hi-fi system was by describing how water in a glass placed on top of the loudspeaker would splash out when the system played loudly. This person held the mistaken impression that the ability to make the water fly from the glass was an indicator of the system's power and quality.

But, his description told me immediately that this loudspeaker system was of *poor* quality. That's because low-quality speakers have thin, vibration-prone cabinets that color the sound. The more a speaker cabinet vibrates, the worse it is. Any speaker system that will splash water out of a glass must sound dreadful. We're about to see why.

When excited by the sound from the driver, primarily the woofers, the enclosure resonates at its natural resonant frequencies. Some of the woofer's back-and-forth motion is imparted to the cabinet, either mechanically through the coupling of the driver to the enclosure, or acoustically by the sound inside the enclosure. This causes the enclosure panels to move back and forth, producing sound. The enclosure thus becomes an acoustic source: we hear music not only from the drivers, but also from the enclosure.

Enclosure vibrations produce sound largely over one or more narrow bands of frequencies centered on the panel's resonant frequencies. The loudspeaker has greater acoustical output at those frequencies. Consequently, cabinet resonances change the timbres of instruments and voices. With a recording of double bass or piano, you can easily hear cabinet resonances as changes in timbre with certain notes. Some musical instruments exhibit a similar phenomenon, producing "wolf tones."

Enclosure resonances not only color the sound spectrally (changing instrumental timbre), they smear the time relationships in music. The enclosure stores energy and releases that energy slowly over time. When the next note is sounded, the cabinet is still producing energy from the previous note. Loudspeakers with severe cabinet resonances produce smeared, blurred bass instead of a taut, quick, clean, articulate foundation for the music. Enclosure vibration also affects the midrange by reducing clarity and smearing transient information.

The acoustic output of a vibrating surface such as a loudspeaker enclosure panel is a function of the excursion of the panel and the panel's surface area.

Because loudspeaker cabinet panels are relatively large, it doesn't take much back-and-forth motion to produce sound. A large panel excited enough that you can feel it vibrating when you play music will color the sound of that loudspeaker. That's why smaller speakers often sound better than similarly priced but larger loudspeakers; it's much harder (and more expensive) to keep a large cabinet from vibrating.

Ideally, the enclosure should create a completely inert platform for the drivers. When the loudspeaker is driven by a high-level transient signal, all of that transient's energy should be converted instantly to acoustic energy by the drive units. The enclosure should be stiff so that it's not easily put in motion, and also damped so that any energy that is imparted to the enclosure by the drive units' motions is quickly suppressed.

Cabinet vibration is why the same loudspeaker sounds different when placed on different stands, or with different coupling materials between the speaker and stand. Cones, isolation feet, and other such accessories all cause the speaker cabinet to have a different resonant signature. I noted with interest a paper, presented at an Audio Engineering Society convention, that explored how much cabinet resonance was audible. Before the researchers could begin the experiment, however, they discovered that the material on which the loudspeaker was resting had an enormous effect on how the cabinet vibrated. For consistency in the experiment, they hung each speaker under test in a rope harness.

Enclosure resonances can be measured and displayed in a way that lets you intuitively see how bad the resonances are, and at what frequencies they occur. (This technique is discussed in [Chapter 16](#), under "Loudspeaker Measurements.") An easier way of judging an enclosure's inertness is the "knuckle-rap" test. Simply knock on the enclosure and listen to the resulting sound. An enclosure relatively free of resonances will produce a dull thud; a poorly damped enclosure will generate a hollow, ringing tone. You can also drive the loudspeaker with a swept sinewave (a special test signal found on some test CDs consisting of a pure tone that gradually and continuously increases in frequency) while resting your hand on the enclosure. When the test signal's frequency momentarily matches the enclosure's resonant frequency, the enclosure will vibrate and you will hear a change in the tone's sound. The less the enclosure vibrates and the fewer the frequencies that cause the enclosure to vibrate, the better.



Fig.6-20 Some loudspeakers feature extraordinary construction designed to prevent the enclosure from vibrating. (Courtesy Magico)

Designers reduce enclosure resonance by constructing cabinets from thick, vibration-resistant material. Generally, the thicker the material, the better. Most loudspeakers are made of $\frac{3}{4}$ "-inch thick *medium density fiberboard* (MDF). MDF that's 1" thick is better; some manufacturers use $\frac{3}{4}$ " on the side panels and top,

and 1" or 2" MDF for the front baffle, which is more prone to vibration. Exotic materials and construction techniques are also used to combat cabinet resonances. These include making the front baffle from Corian (a synthetic material used in kitchen countertops), aluminum, or carbon fiber. One loudspeaker designer molded a baffle from a mixture of concrete and special polymers; it weighed more than 100 pounds.

Cross-braces inside the enclosure reduce the area of unsupported panels and make the cabinet more rigid. Some braces are called "figure-8" braces because their four large holes form a figure-8 pattern. These braces are strategically placed for maximum effectiveness. The importance of cabinet bracing can be seen in [Fig.6-20](#); note the elaborate aluminum internal structure, whose sole purpose is to prevent enclosure vibration.

Another technique for building a stiff and well-damped enclosure is shown in [Fig.6-21](#), which shows an enclosure in cross section. The enclosure is built from two skins of fiberglass or carbon fiber (the latter in more expensive models), with the space between the skins filled with a dense epoxy resin. The skins provide the stiffness, the epoxy the damping. This way of building an enclosure has the advantage of requiring no joints or seams, while making it possible to use compound curves, which further reduce cabinet diffraction. This innovative "molded composite" technique was developed by Andy Payor, of Rockport Technologies, for use in his company's upper-end models.

The modern trend toward building inert loudspeaker enclosures was started by Wilson Audio in 1986, with the WATT loudspeaker. This small, two-way speaker employed lead panels fitted in hollowed-out sections of the cabinet walls. Today's Wilson loudspeakers are made from a combination of materials Wilson calls "X" and "M", both of which are resin-based ceramics that can be machined like metal.

Much of the improvement in loudspeaker systems over the past 20 years can be attributed to a greater understanding of the importance of reducing cabinet vibration



Fig.6-21 Cabinet cross section of Rockport Technologies' "molded composite" enclosure

Enclosure Shapes

The enclosure can also degrade a loudspeaker's performance by creating *diffraction* from the cabinet edges, grille frame, and even the drivers' mounting bolts. Diffraction is a re-radiation of energy that happens when the sound encounters a discontinuity in the cabinet, such as at the enclosure edge. Diffracted energy combines with the direct sound from the drivers to produce ripples (i.e., colorations) in the frequency response. Rounded baffles, recessed drivers and mounting bolts, and low-profile grille frames all help to reduce diffraction. The midrange enclosure shown in [Fig.6-17](#) (B&W's Nautilus technology) is an example of diffraction reduction taken to an extreme.

Some loudspeaker employ sloped baffles to align the drivers in time. This slope aligns the acoustic centers of all the drivers so that their outputs arrive at the listener at the same time (see [Fig.6-22](#)). Although many loudspeaker manufacturers use the words "time-aligned," the term is a trademark owned by loudspeaker designer Ed Long.




Fig.6-22 A sloped baffle causes the sound from all of the speaker's drivers to reach the listener simultaneously. (Courtesy Thiel Audio)

Crossovers

A loudspeaker *crossover* is an electronic circuit inside the loudspeaker that separates the audioband into different ranges of frequency, and sends each to the appropriate drive unit: bass to the woofer, midband frequencies to the midrange, and treble to the tweeter (in a three-way loudspeaker).

A crossover is made up of capacitors, inductors, and resistors. These elements selectively filter the full-bandwidth signal driving the speaker, creating the appropriate filter characteristics for the particular drivers used in that speaker. A crossover is usually mounted on the loudspeaker's inside rear panel or bottom panel.

A crossover is described by its cutoff frequency and slope. The cutoff frequency is the frequency at which the transition from one drive unit to the next occurs—between the woofer and midrange, for example. The crossover's slope refers to the rolloff's steepness. A slope's steepness describes how rapidly the response is attenuated above or below the cutoff frequency. For example, a first-order crossover has a slope of 6dB/octave, meaning that the signal to the drive unit is halved (a reduction of 6dB) one octave above the cutoff frequency. If the woofer crossover circuit produces a cutoff frequency of 1kHz, the signal will be rolled off (attenuated, or reduced in level) by 6dB one octave higher, at 2kHz. In other words, the woofer will receive energy at 2kHz, but that energy will be reduced in level by 6dB. A first-order filter producing a 6dB/octave slope is the slowest and most gentle rolloff used.

The next-steeper filter is the second-order crossover, which produces a rolloff of 12dB/octave. Using the preceding example, a woofer crossed over at 1kHz would still receive energy at 2kHz, but that energy would be reduced by 12dB (one quarter the amplitude) at 2kHz. A third-order crossover has a slope of 18dB/octave, and a fourth-order crossover produces the very steep slope of 24dB/octave. Using the previous example, the fourth-order filter would still pass 2kHz to the woofer, but the amplitude would be down by 24dB (1/16th the amplitude). [Fig.6-23](#) compares crossover slopes.

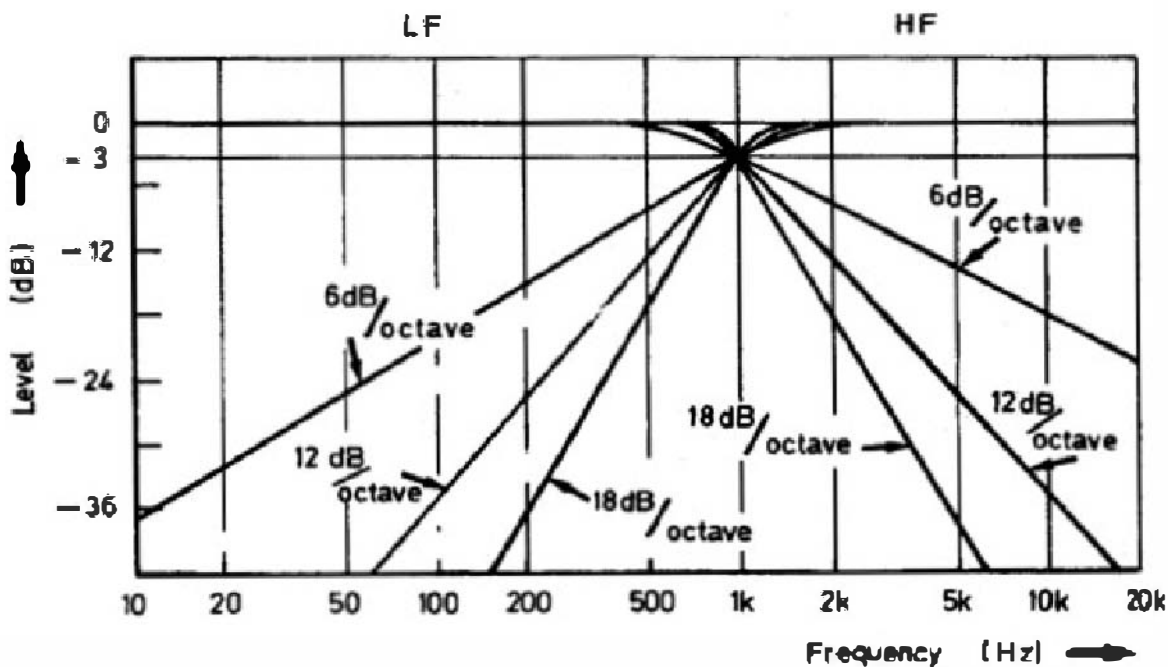


Fig.6-23 A comparison of crossover slopes (Courtesy Martin Colloms, *High Performance Loudspeakers*)

A typical crossover frequency for a two-way loudspeaker would be somewhere between 1 and 2.5kHz. A three-way system may have crossover frequencies of 800Hz and 3kHz: The woofer reproduces frequencies up to 800Hz, the midrange driver handles the band between 800Hz and 3kHz, and the tweeter reproduces frequencies above 3kHz.

Note that the actual acoustic crossover slope—the drive unit’s acoustical output—may be different from the electrical slope produced by the crossover. If a drive unit is operated close to its own rolloff frequency, this inherent rolloff is added to the electrical rolloff. For example, a woofer rolling off naturally at 1kHz at 6dB/octave will produce an acoustic rolloff of 12dB/octave when crossed over at 1kHz with an electrical 6dB/octave crossover.

The crossover slope is a major design decision. A shallow first-order slope produces the least phase error in the crossover region between the two drive units. The two drive units operate in phase at both the crossover point and outside the crossover region. That is, the sounds produced by the various drivers are in step with each other in time. A fourth-order filter, by contrast, can cause a phase lag between drivers; the individual drivers’ output aren’t produced at exactly the same time. Loudspeakers with second-order crossovers usually have

their tweeters wired in inverse polarity to correct for the crossover's phase lag and provide correct acoustic polarity.

Simple first-order crossovers have inherently better time behavior than more complex crossovers. First-order crossovers, however, produce considerable overlap between drive units. There is a relatively wide band over which both the woofer and the midrange driver produce the same frequency, and over which the midrange and tweeter overlap. This can cause dips in the frequency response if the listener's ears aren't exactly the same distance from each drive unit. The outputs of the two drivers can combine, causing destructive interference and irregular amplitude response. The listener should also sit farther away from a speaker with first-order crossovers, to give the sounds from the individual drive units more time to integrate.

In addition, first-order crossovers can allow the tweeter to be overdriven by midrange frequencies in the octaves below the nominal cutoff frequency. A tweeter's excursion quadruples with each halving of frequency, moving four times farther at 1kHz than at 2kHz, given the same drive level. At 500Hz its motion quadruples again, if it isn't rolled off by the tweeter's fundamental resonance. Although a 6dB/octave crossover slope halves the drive signal with each halving of frequency, that's not enough to offset the tweeter's natural quadrupling of excursion as the frequency is halved. The result can be reduced power handling by the system as the drivers are driven to their excursion limits at a lower sound-pressure level. A tweeter driven by excessive midband energy can also introduce distortion, which is manifested as a hard and metallic sound. Proponents of high-order crossovers suggest that any phase advantages of first-order slopes are more than offset by the unwanted overlap between drivers, and the concomitant potential for driver overload. A fourth-order crossover provides sharper delineation of the frequency ranges of each pair of drivers.

A factor in choosing crossover points is the individual drivers' directivity—their dispersion of sound with frequency. As the frequency a driver is producing increases, the dispersion pattern becomes narrower. In other words, bass is omnidirectional and high frequencies tend to beam. Driver size and crossover points are chosen so that there is no discontinuity in directivity between drivers. This is why you don't see a 10" woofer mated to a 1" tweeter—the discontinuity between the two drivers' dispersion patterns would be audible.

Some loudspeakers can be connected to a power amp with two runs of loudspeaker cable, a technique called *bi-wiring*. A loudspeaker with provision for bi-wiring has separate high-pass (tweeter) and low-pass (woofer) crossover sections, along with two pairs of input terminals. You should always take

advantage of this feature by running two pairs of loudspeaker cables: Bi-wiring can significantly improve a loudspeaker's sound. (Bi-wiring is discussed in more detail in [Chapter 10](#))

Crossovers vary greatly in the quality of parts used. Budget loudspeakers will likely use iron-core inductors rather than the preferable air-core types. Air-core inductors are larger and more expensive than iron-core coils, but are immune to a phenomenon called saturation, which introduces distortion. Similarly, lower-cost loudspeakers use electrolytic capacitors rather than the more expensive plastic-film types. Electrolytic capacitors are sometimes bypassed with small-value polypropylene or other similar high-quality capacitors. Bypassing is putting a smaller-value, higher-quality capacitor in parallel with the larger-value capacitor. Some of the sonic virtues of high-quality capacitors are retained, accompanied by the cost and size advantages of an electrolytic capacitor. Many ultra-high-end loudspeakers use very expensive gold- and silver-foil capacitors that can cost as much as \$300 each. Finally, some loudspeakers are wired internally with high-end cable rather than generic wire.

Digital Loudspeakers

The capacitors, inductors, and resistors that make up a loudspeaker crossover are far from perfect. Not only do they exhibit variations in value that affect performance, these components split up the audioband in a relatively crude way. As mentioned, higher-order crossovers produce better division of the audioband, but their phase response suffers. Moreover, a significant amount of amplifier power is wasted in the crossover.

The widespread availability of digital audio sources provides an opportunity to remove traditional crossovers from the signal path. If we can divide the audioband digitally—that is, by performing mathematical computations on the audio signal—we can create just about any crossover characteristics we want, with none of the problems of capacitors, inductors, and resistors.

These *digital loudspeakers* accept a digital input signal and implement the crossover in the digital domain with *digital signal-processing* (DSP) chips. Instead of subjecting the audio signal to resistors, capacitors, and inductors, DSP crossovers separate the audioband by performing mathematical processing on the digital audio data. DSP crossovers can have perfect time behavior, any slope and frequency the designer wants without regard for component limitations or tolerances, and apply equalization to the individual drive units.

[Fig.6-24](#) is a side view of a combination pictorial/block diagram of a digital loudspeaker. The speaker accepts a digital input signal from a disc player, music server, or other digital source. DSP chips inside the speaker split the audioband into bass, midrange, and treble. These chips can also equalize and delay signals to produce nearly perfect acoustic behavior at the drive units' output. Each of the three digital signals (bass, midrange, treble) is then converted into an analog signal with its own digital-to-analog converter (DAC). The DAC output feeds a power amplifier specially designed to power the particular drive unit used in the digital speaker. This aspect of a digital speaker confers a large advantage: the power amplifiers are designed to drive a known load. The power amplifiers amplify an analog signal, just like a conventional power amp, then drive conventional cone speakers.

Digital loudspeakers offer the designer the ability to equalize the signal driving the woofer to achieve deeper bass extension for a given enclosure volume than is possible in a conventional loudspeaker. Earlier in this chapter we saw how a woofer in an enclosure begins rolling off just below its resonant frequency, with

the resonant frequency and rolloff slope determined by the acoustical loading. In a digital loudspeaker, the signal driving the woofer can be boosted to compensate for this rolloff, extending the low-frequency response. With this technique, it's possible to achieve amazingly deep bass extension from a moderately sized enclosure.

In addition to employing virtually perfect crossovers, digital loudspeakers also simplify your system. You need only a digital source and a pair of loudspeakers for a complete music system. This eliminates monoblock power amplifiers on the floor and long runs of garden-hose speaker cable. Moreover, the digital nature of the system provides for remote control of the speakers. You can adjust the tonal balance, correct for different listening heights, and even delay one loudspeaker if your room doesn't permit symmetrical speaker placement. The digital speaker will also have a volume readout, which makes setting the correct playback level for different music more convenient. When considering the price of digital loudspeakers, keep in mind that you don't need to buy an outboard digital-to-analog converter or power amplifiers.

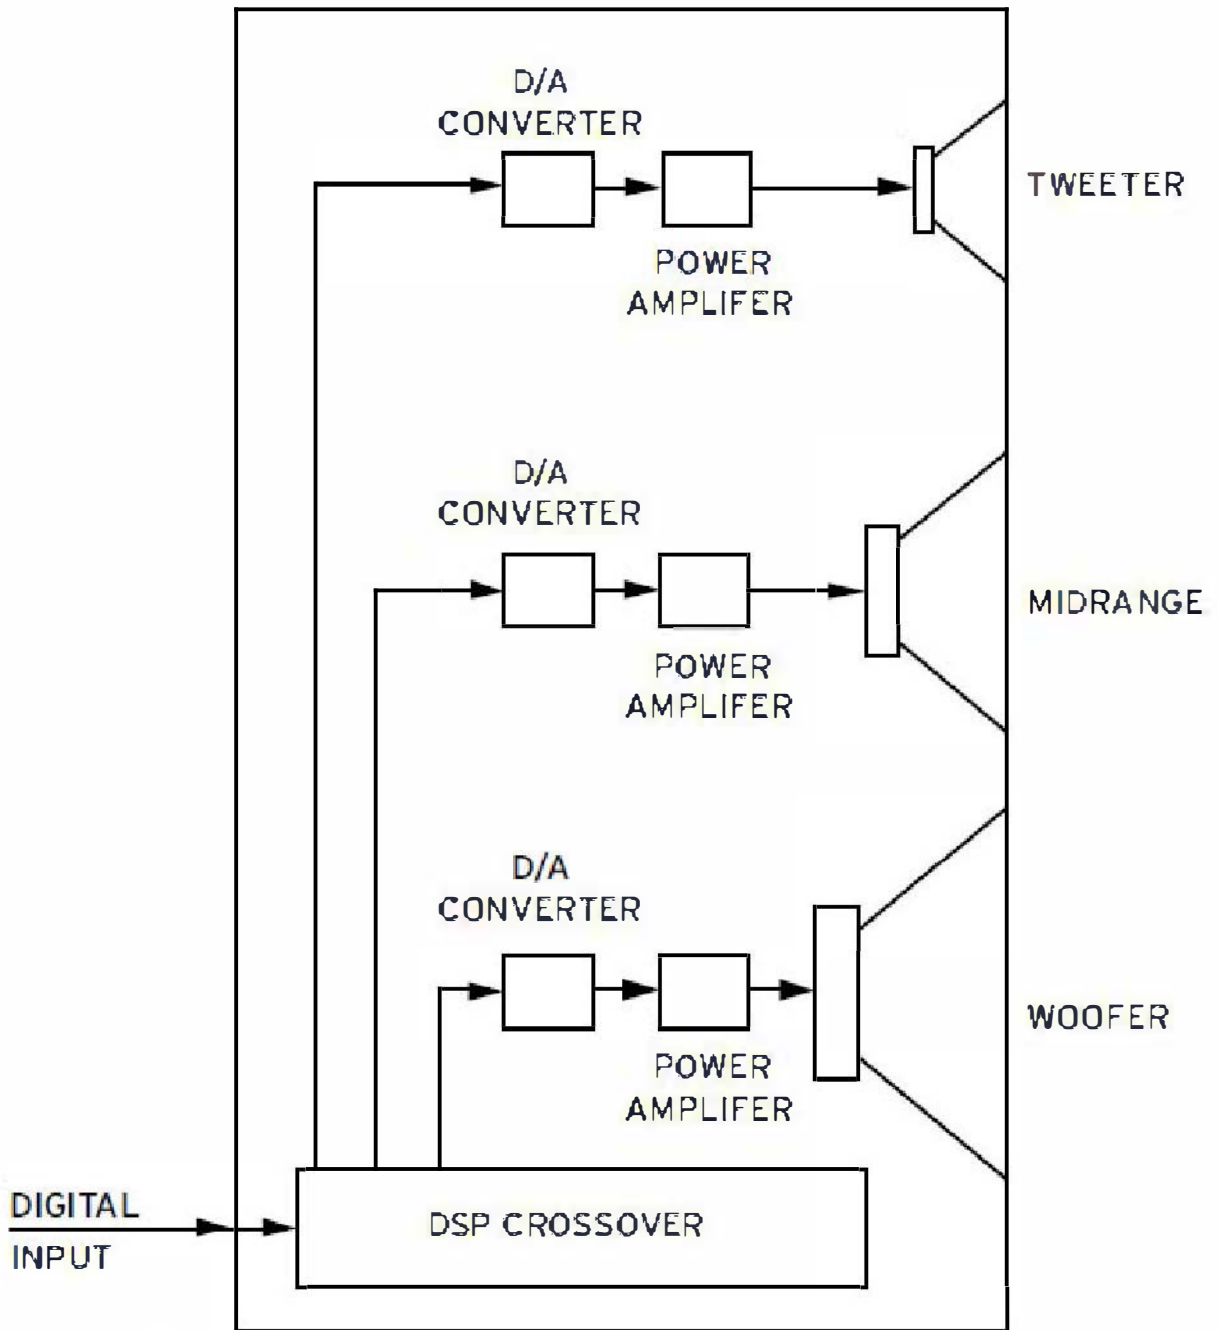


Fig.6-24 A digital loudspeaker incorporates separate digital-to-analog converters and power amplifiers for each driver.

A digital speaker also “knows” the sound-pressure level the loudspeaker is producing and can add a continuously variable loudness compensation. You may have seen the Loudness button on some audio products, which engages a circuit that boosts bass and treble at low volumes to compensate for the ear’s reduced sensitivity to bass and treble at low volume. (See [Appendix A](#) for a complete explanation.) Although loudness compensation is, in theory, a good idea, loudness controls are eschewed by audiophiles because they are crude, and effective at only one precise volume setting. A digital loudspeaker, by contrast, can apply exactly the right amount of loudness compensation at any volume setting. (The correct amount of bass and treble boost for a given sound-pressure level is precisely known from decades of psychoacoustic research.)

Another advantage of a loudspeaker under DSP control is the ability to protect the drivers from damage; the DSP knows the driver’s limitations and can impose a maximum excursion limit.

Digital loudspeakers were pioneered by Meridian Audio, which has been producing them since 1993. Meridian’s first models accepted the 44.1kHz 16-bit datastream from a CD transport, but newer models have been upgraded to accept high-resolution digital audio with sampling rates of up to 192kHz and word lengths of up to 24 bits. Newer models also feature vastly more powerful DSP engines that deliver better performance than earlier models.

DSP Loudspeakers

Not to be confused with digital loudspeakers, *DSP loudspeakers* employ digital signal processing (DSP) chips to produce flat frequency response in a room. By contrast, digital loudspeakers use DSP to achieve flat response from the loudspeaker, but ignore the room's effect on frequency response.

With a DSP loudspeaker, you measure the frequency response of the loudspeaker/room combination at the listening seat. The DSP uses this measurement to create a set of filters that are the inverse of the frequency response. That is, a measured peak of a certain bandwidth at, for example, 70Hz results in a dip in the filter at that frequency of precisely the same amplitude and bandwidth as the measured peak. The filters address hundreds of such frequency-response deviations. In essence, the filters introduce a frequency-response error into the signal driving the loudspeakers so that the audio signal, when modified by the loudspeaker/room combination, has a flat frequency response. DSP speakers often give you the opportunity to create custom frequency-response curves in addition to the flat curve (which often sounds too bright). Note that if you move your speakers or listening position, then you must re-measure the response and let the DSP loudspeaker create a new set of filters.

In practice, the DSP loudspeaker is supplied with an outboard processing unit that fits between your preamplifier and power amplifier. The processor converts analog input signals to digital, applies the correction algorithm, and converts the digital signal back to analog for output to your power amplifier. In addition to correcting the frequency response, the processor can time-align the arrival of the drivers' outputs at the listening position to deliver flat phase response. It often also contains a crossover with level adjustments.

In essence, the DSP speaker is a conventional speaker that has been designed specifically to be used with the supplied DSP unit.

Subwoofers

A subwoofer is a loudspeaker that produces low frequencies that augment and extend the bass output of a full-range loudspeaker system. The term *subwoofer* is grossly misused to describe any low-frequency driver system enclosed in a separate cabinet. But *subwoofer* actually means “below the woofer,” and should be reserved for those products that extend bass response to below 20Hz. A low-frequency driver in an enclosure with an output extending to 40Hz and used with small satellite speakers is more properly called a woofer.

You’ll also see full-range speakers with a built-in “subwoofer” powered by its own amplifier. Most of these products actually employ woofers that are simply driven by an integral power amplifier. Such a design relieves your main amplifier of the burden of driving the woofer, but requires that the speakers be plugged into an AC outlet.

Subwoofers come in two varieties: passive and active. A passive subwoofer is just a woofer or woofers in an enclosure that must be driven by an external amplifier. In one variation of the passive subwoofer, the same stereo amplifier driving the main speakers also powers the subwoofer. In this least desirable method of connecting a subwoofer, the full-range output from a power amp is input to the subwoofer, and a crossover in the subwoofer removes low frequencies from the signal and outputs the filtered signal to the main loudspeakers. This technique puts an additional crossover in the signal path, to filter out bass from the speaker-level signals. Some subwoofers are designed to extend the bass response of the main speakers without filtering bass from the signal driving them. Although this technique doesn’t add a crossover to the signal path, it doesn’t increase the power handling of the main speakers. Keeping low bass out of your main speakers confers many advantages, including a much cleaner midrange.

A better way of driving the subwoofer is with an *electronic crossover* and separate power amplifier. This method separates the bass from the signal driving the main loudspeakers at line level, which is much less harmful to the signal than speaker-level filtering. Moreover, adding a separate power amp for the subwoofer greatly increases the system’s dynamic range and frees the main-speaker amplifier from the burden of driving the subwoofer. Adding a line-level crossover and power amp turns the passive subwoofer into an active subwoofer, and also makes the system bi-amplified.

A self-contained active subwoofer combines a subwoofer with a line-level crossover and power amplifier in one cabinet, eliminating the need for separate boxes and amplifiers (Figs.6-25 and 6-26). Such a subwoofer has line-level inputs (which are fed from the preamplifier), line-level outputs (which drive the power amp), and a volume control for the subwoofer level. The line-level output is filtered, removing low frequencies from the signal sent to the amplifiers driving the main loudspeakers. This crossover frequency is adjustable to allow you to select the frequency that provides the best integration with the main speakers.

A useful subwoofer feature is an automatic equalization circuit that measures the frequency response of your system in your room and tailors the subwoofer's output to provide flatter response. You connect a supplied calibration microphone to the subwoofer, push a button, and the subwoofer emits a series of tones that are picked up by the microphone and analyzed by a circuit in the subwoofer. The subwoofer then equalizes its output, applying a boost to certain frequencies and a cut to others so that the resultant output is as flat as possible.

Adding an actively powered subwoofer to your system can greatly increase its dynamic range, bass extension, midrange clarity, and ability to play louder without strain. The additional amplifier power and low-frequency driver allow the system to reproduce musical peaks at higher levels. Moreover, removing low frequencies from the signal driving the main loudspeakers lets them play louder because they don't have to reproduce low frequencies. The midrange often becomes clearer because the woofer cone isn't furiously moving back and forth, trying to reproduce low bass. This improvement in a small speaker's performance can be dramatic. The small woofer's excursion is no longer a limiting factor in how loudly the system will play when low frequencies are filtered from the signal driving it. The midrange is much cleaner, and the overall system sounds like a large, full-range speaker.



Fig.6-25 An active subwoofer with integral crossover, designed for music reproduction. (Courtesy JL Audio)



Fig.6-26 An active subwoofer with integral crossover provides a variety of controls for optimizing the subwoofer's performance in your room. (Courtesy JL Audio)

Now for the bad news. More often than not, subwoofers degrade a playback system's musical performance. Either the subwoofer is poorly engineered (many are), set up incorrectly, or, as is increasingly common, it's designed to reproduce explosions in a home-theater system, not resolve musical subtleties.

Let's first look at the theoretical problems of subwoofers. First, most subwoofers—passive or active—add electronics to the signal path. The active subwoofer's internal crossover may not be of the highest quality. Even well-executed crossovers can still degrade the purity of very-high-quality source components, preamplifiers, and power amps. This drawback can be avoided by running the main loudspeakers full-range (no rolloff), but you then lose the dynamic advantages and additional midrange clarity conferred by keeping low frequencies out of the main speakers.

Second, the subwoofer's bass quality may be poor. The subwoofer may move lots of air and provide deep extension, but a poorly designed subwoofer often adds a booming thumpiness to the low end. Rather than increase your ability to hear what's going on in the bass, a subwoofer often obscures musical information.

Third, a subwoofer can fail to integrate musically with the main loudspeakers. Very low frequencies reproduced by the subwoofer can sound different from the midbass produced by the main speakers. The result is an extremely distracting discontinuity in the musical fabric. This discontinuity is manifested as a change in the sound of, for example, acoustic double bass in different registers. Ascending and descending bass lines should flow past the crossover point with no perceptible change in timbre or dynamics.

Another factor that can make integrating a subwoofer difficult is matching a slow subwoofer to taut, lean, articulate main speakers. Put another way, the sound from an underdamped subwoofer won't integrate very well with that from a pair of overdamped speakers.

Fourth, subwoofers often trade tight control, pitch resolution, and lack of overhang for greater sensitivity or deeper extension. This is particularly true of subwoofers designed for home theater. Consequently, many subwoofers sound bloated, slow, and lacking in detail.

Finally, a subwoofer can fill the listening room with lots of low-frequency energy, exciting *room-resonance modes* that may not have been that bothersome without the subwoofer. Placement is therefore crucial—you can't put a subwoofer just anywhere and expect musical results. This problem of room-mode excitation

can be ameliorated by two (or more) subwoofers; each subwoofer will excite different room modes, substantially smoothing out the room's low-frequency response.

All of these problems are exacerbated by most people's tendency to set subwoofer levels way too high. The reasoning is that if you pay good money for something, you want to hear what it does. But if you're aware of the subwoofer's presence in the sound, either its level is set too high, it isn't positioned correctly, or the subwoofer has been poorly designed. The highest compliment one can pay a subwoofer is that its contribution can't be directly heard. It should blend seamlessly into the musical fabric, not call attention to itself.

Having said all that, it is possible to greatly improve a system's performance by adding a subwoofer. In fact, there are many compelling arguments that a subwoofer/satellite system has the potential to outperform a similarly priced full-range loudspeaker system. These arguments include:

- The satellite speakers can be positioned for the best soundstage without regard for how that position affects the bass response.
- The subwoofers can be positioned for the best bass integration in the room without regard for soundstaging.
- It is much more cost-effective to build a subwoofer enclosure than it is to build the large enclosure of a full-range loudspeaker (to accommodate large woofers), which may have expensive cabinet construction and wood or paint finishes. The subwoofer is less likely to be a prominent part of the home decor, and thus doesn't require lavish finish quality.
- The subwoofer/satellite system gives you a large measure of control over the bass performance, including the amount of bass and how that bass integrates with the room. A full-range loudspeaker offers no such control.

The counterargument suggests that a loudspeaker system should be designed and engineered with a single vision, not with a piecemeal approach by different designers. My experience suggests that a subwoofer/satellite system can outperform a similarly priced full-range loudspeaker system, but with two big caveats. The first is that the subwoofer must be of exceptional quality and be designed for musical performance, not booming home-theater effects. Only a tiny handful of the hundreds of subwoofers on the market fall into this category. Second, achieving seamless integration between the subwoofer and satellites and

realizing smooth bass response requires skill and patience. The subwoofer/satellite system is a better choice for someone who enjoys the technical side of audio rather than for the music lover who just wants good sound without having to ascend the learning curve.

A step-by-step guide to getting the best sound from a subwoofer is included in [Chapter 14](#) (“System Setup Secrets: Loudspeaker Placement and Room Acoustics”).

Subwoofer Technical Overview

Designers of subwoofers face the same choices of low-frequency loading that designers of full-range speakers do. The woofer or woofers can be loaded into an infinite baffle (sealed box), reflex enclosure (ported), or a transmission line. Transmission-line subwoofers are rare because of the large cabinet required for such a design. Instead, the sealed-box and reflex are the most popular. Some subwoofers use servo drive (explained earlier in this chapter).

The simplest subwoofer is a single driver—usually a 10” or 12” cone—in a reflex enclosure. The choice of a reflex design keeps the cabinet at a moderate size and maintains reasonable sensitivity. A larger enclosure extends the subwoofer’s bass response, but is more expensive to manufacture, particularly if the enclosure is properly constructed (i.e., well braced).

The sealed-box subwoofer is less common due to its inherently higher cutoff frequency. A sealed subwoofer would have to be enormous to achieve sub-20Hz extension; most sealed-box units are actually woofers, not subwoofers. They trade bass extension for bass of higher-quality: better control, faster transient response, and precise pitch definition, for example.

Some small sealed subwoofers can, however, achieve sub-20Hz extension because the signal driving them has been equalized to produce flat response. The equalization applies a boost to the low bass that counteracts the woofer’s natural rolloff, producing flat response far below the woofer’s natural cutoff frequency. Equalization puts enormous stress on the woofer cone because it forces the woofer to handle a large amount of power and have high excursions at very low frequencies. Consequently, a woofer designed for use in an equalized system will have a large magnet and huge voice coil. If an equalized subwoofer uses a separate, rather than integral, power amplifier and equalization circuit, be aware that the amplifier will work correctly only with the subwoofer for which it was designed. Similarly, a subwoofer designed to be driven by an equalized signal will work correctly only when driven by its matching amplifier.

Some subwoofers are mounted in bandpass enclosures that limit the system’s bandwidth with both electrical and acoustical filtering. In a bandpass enclosure, the woofer drives a chamber that is vented through a port or slot. The woofer is often loaded in a smaller, rear chamber. This version is called a second-order bandpass alignment. Porting the smaller chamber results in a fourth-order alignment, producing steeper filter slopes. Advantages of bandpass enclosures

include greater low-frequency extension, lower distortion, greater sensitivity, and higher power handling.

Subwoofers usually offer a choice of crossover frequencies, either at fixed points (typically 80, 120, and 180Hz) or with a continuously variable knob. A higher crossover frequency would be selected for a minimonitor than for a full-range loudspeaker. Generally, the lower the crossover frequency, the better; the main speaker's bass is probably of higher quality than the subwoofer's, and a low crossover frequency moves any crossover discontinuity lower in frequency, where it will be less audible. In addition, a low crossover frequency ensures that you won't be able to locate the sound source of the low bass. A subwoofer reproducing frequencies above 100Hz can be localized—i.e., the location of the source of the bass can be detected—which is musically distracting. A crossover frequency that is too low will, however, burden small speakers with excessive bass and reduce the system's power handling and maximum listening level. I've used a subwoofer with a massive full-range loudspeaker system, setting the crossover frequency at just 30Hz. The subwoofer provided just a little more extension and power in the bottom octave. I've also successfully used a subwoofer with a crossover setting as high as 100Hz.

Another variable in subwoofer crossovers is the slope. Most use fourth-order (24dB/octave) filters. Ideally, the crossover frequency and slope would be tailored to the particular speaker used with the subwoofer. But because the subwoofer manufacturer doesn't know which speakers will be used with the subwoofer, these parameters are compromised for good performance with a variety of loudspeakers.

Not all subwoofer crossovers are built into the subwoofer. It's possible to buy an outboard crossover that offers more flexibility than a subwoofer's integral crossover (Fig.6-27). This crossover offers a control marked Damping. To understand this adjustment, think about the main speakers and the subwoofer reproducing the same frequency at the crossover point. How those two signals sum acoustically can be unpredictable. The Damping adjustment allows you to fine-tune this transition by applying a gentle boost or cut at the crossover frequency. This adjustment can be extremely useful in integrating the output from the subwoofer and the main speakers.



Fig.6-27 An outboard line-level subwoofer crossover (Courtesy JL Audio)

Loudspeaker Stands

For best performance, small loudspeakers should be mounted on stands. In fact, the stands' quality can greatly affect the reproduced sound. Flimsy, lightweight stands should be avoided in favor of solid, rigid models. The stands should include spikes on their bottoms, to better couple stand and speaker to the floor. Some speaker stands can be filled with sand or lead shot for mass loading, to make them more inert and thus less prone to vibrate. A great loudspeaker on a poor stand will suffer significantly degraded performance. Plan to spend several hundred dollars on stands. When comparing a floorstanding loudspeaker to one requiring stand-mounting, include the cost of the stands in your budgeting.

The interface between loudspeaker and stand also deserves attention. Spikes, cones, and other isolation devices (see [Chapter 15](#)) can allow the loudspeaker to perform at its best. In a series of experiments attempting to quantify the effects of different materials installed between stand and speaker, an accelerometer was attached to the speaker's side panel and its vibration measured with the different interface materials. The speaker's resonant "signature" was changed considerably by the stand, and again by the material or device between stand and speaker. Most effective was a small ball of Bostik Blu Tack placed at each corner of the speaker stand. This very effective yet inexpensive interface is a sticky, gum-like material available at hardware stores. Best of all, a lifetime supply costs about \$2.

When setting up speaker stands, be sure that each stand's three or four spikes each contacts the floor, and that the stand can't rock back and forth. You can also improve the coupling of speaker to stand, and stand to floor, by putting a heavy object on top of the speaker. Commercial products, made from a heavy granular material in a canvas bag, are available for this purpose. The bags also help damp enclosure resonances.

7 Digital Source Components, Part One: Disc Players, DACs, and Interfaces

Introduction

A digital source component is any product or combination of products that plays music from a digital medium. A digital front-end can be as simple as a CD player or as complex as a combination of a *multiformat disc player*, Internet-connected *music server*, *wireless streamer*, outboard *digital-to-analog converter*, and separate *master clock*. Other digital source components include *CD/SACD players*, *Blu-ray Disc players*, and *universal players* that support all disc formats. Whatever the configuration, your system's digital front-end is a vital link in the audio reproduction chain, and the source you'll probably spend the most time listening to.

No area of audio has experienced greater technological change over the past few years than the digital arena. We've seen a revolution in how we access, store, and play digitally encoded music. Just a few years ago, digital audio was limited to listening to an optical disc (CD or SACD). Today, we can store our entire music libraries on a hard-disk drive, access those music libraries via a phone or tablet app, wirelessly stream audio from a subscription-based music service to multiple systems in the home, manage a large music collection in a way that makes it easy to find just what we're looking for, and explore an artist or genre with virtually unlimited and up-to-date "liner notes" via the Internet.

The idea of purchasing a physical format such as a disc, removing that disc from a package, and putting it into a player now seems to some music lovers like a quaint anachronism. Nonetheless, CD and SACD remain important formats for the audiophile. SACD, in particular, offers exceptional sound quality and a

surprisingly large catalog of titles, particularly of classical music. Optical discs may be on their way out, but they still have a few more years of useful service left.

New digital technologies are blurring the lines between traditional component categories (integrated amplifiers with DACs, for example), as well as introducing entirely new classes of products. Many of these revolutionary new products eschew the rectangular boxes of yesterday's components for small, sleek shapes that look nothing like yesterday's audio products.

Because digital audio spans such a wide spectrum of products and technologies, I've split the topic into two chapters. This chapter covers how to choose a digital source, what to listen for, disc players, DACs, and digital interfaces. The following chapter delves into music servers, computer-based audio, and *networked systems*.

How to Choose a Digital Source-Overview

The first consideration in choosing digital source components is deciding which formats to invest in. For many music lovers, particularly those who don't want to boot up a computer in order to listen to music, a CD player will suffice. Others will opt for a CD/SACD player to take advantage of SACD's superior sound quality. Although SACD is generally considered a "failure" in the mainstream market, it remains a robust format for audiophiles, particularly those who enjoy classical music. The SACD catalog abounds with outstanding performances in terrific sound. Moreover, SACD is the format for high-resolution, multichannel digital audio; no other medium, whether packaged or downloaded, makes available nearly as many titles of high-res, multichannel recordings as does SACD. Many record companies still regularly release new titles in multichannel SACD, and some release *only* hybrid CD/SACD discs that are compatible with CD players. If you enjoy classical music in multichannel sound, SACD is the format of choice.

The DVD-Audio format, once proposed as the successor to the Compact Disc, is a dead format. Still, some universal disc players offer DVD-A compatibility. The Blu-ray Disc (BD) format has the technical capacity to serve as a high-resolution, multichannel, music-only format, but so far the world's record companies have largely ignored its potential as a music carrier. Some specialty record labels have embraced the BD and offer a range of music titles, some accompanied by high-definition video of the artists in concert or in a studio performance. The BD format has much to offer the audiophile; it can deliver concerts and musical performances in high-definition video, along with high-resolution multichannel audio that has the same quality as the studio master. Nonetheless, the major record companies see a future based not on physical formats but on downloads and, increasingly, a subscription-based music-distribution model (more on this in the following chapter).

Even if your system includes a disc player, you'll probably also want a music server. This term encompasses a wide range of approaches, from a *turnkey server* to a *do-it-yourself*, computer-based system. We'll explore music servers in depth in the next chapter, but an overview of their strengths and weaknesses might help point you in the right direction. On the plus side of the ledger, a music server gives you instant access to your music library by showing you your entire collection on a computer monitor or via an app on your phone or tablet. Playing

a particular piece of music is as simple as clicking a mouse button or tapping a touchscreen, allowing you to spend more time listening and less time finding the CD case and loading a disc in a player. A music server also allows you to program a large number of tracks in advance for uninterrupted listening. Some servers will assemble for you a long playlist based on you telling the server what kind of music you'd like to hear. Living with a music server fundamentally changes your relationship to your music library; I've found myself listening deeper into my collection as more titles are brought to my attention. In addition, music servers open up a vast new world of music via downloads; you can add music to your library without leaving your listening chair. Anyone who has used Apple's iTunes knows the power of a music server. And *downloading* music to a server is the best way to access high-resolution digital audio. Downloads have precipitated the recent explosion in the number of high-res titles; the record companies and giant hardware manufacturers didn't have to spend a decade wrangling over the specifications of a new physical format.

And now, the downside of music servers: Unless you opt for an expensive turnkey system, you must go through the hassle of installing software and hardware, configuring the system, getting it to work correctly, and navigating the learning curve to realize high-quality sound. A do-it-yourself music server with state-of-the-art sound quality requires patience and above-average technical expertise with personal computers. Moreover, some music lovers find that turning on a computer in order to enjoy music is anathema. Those who want the convenience of a server without the hassles can choose a turnkey system that also offers an outstanding user interface, but such systems are expensive, and sometimes prevent you from moving your laboriously created music library to a different hardware platform. We'll consider the differences between turnkey and do-it-yourself servers in the next chapter.

When choosing digital source products, keep in mind that, of all the components in an audio system, the digital front-end is the most likely to be left behind by technology's inexorable progress.

Whatever format(s) you choose for your music system, those digital bits must be converted into an analog signal, and it's the digital-to-analog conversion stage that has the greatest influence on sound quality. Without a great-sounding DAC—whether a separate component or a DAC integral to a disc player, preamplifier, or integrated amplifier—it doesn't matter what formats you've chosen or how sophisticated your music server is. The DAC is where the sonic rubber meets the road, and it will exert an enormous influence on how much you enjoy the music you play back through it.

Before we look at disc players, DACs, and digital interfaces, let's look at how to sonically evaluate digital source components.

What to Listen For

Perhaps more than any components other than loudspeakers, digital sources come in the most flavors—their sonic and musical characteristics vary greatly between brands and models. This variability has its drawbacks (“Which one is right?”), but also offers the music lover the chance to select one that best complements his playback system’s characteristics and suits his musical tastes. The different types of sonic presentations heard from disc players and DACs tend to reflect their designers’ musical priorities. If the designer’s parts budget and/or skill is limited, certain areas of musical reproduction will be poorer than others. The trick is to find the digital source that, in the context of your system, excels in the areas you find musically most important.

Selecting a digital source specifically tailored to the rest of your playback system can sometimes ameliorate some of that playback system’s shortcomings. This is true for all components, but especially so for digital because of their wide variations in sonic signatures. For example, don’t choose a bright-sounding DAC for a system that is already on the bright side of reality. Instead, you may want to select a DAC whose main attribute is a smooth, nonfatiguing treble. Each digital product has its particular strengths and weaknesses. Only by careful auditioning—preferably in your own system—can you choose the product best for you.

To illustrate this, I’ve invented two hypothetical listeners, each with a different system and taste, and two hypothetical DACs—though two disc players could easily be substituted. Although the following discussion could apply to all audio components, it is particularly relevant to digital gear. Not only are there wide variations in sonic characteristics among DACs, but a poor-sounding DAC at the front-end of a superb system will ruin that system’s overall sound.

Listener A likes classical music, particularly early music, baroque, and choral performances. She rarely listens to full-scale orchestral works, and never plays rock, jazz, or pop. Her system uses inexpensive solid-state electronics and somewhat bright-sounding speakers; the combination gives her a detailed, forward treble that is a bit aggressive.

Listener B wouldn’t know a cello from a clarinet, preferring instead electric blues, rock, and pop. He likes to feel the power of a kick drum and bass guitar working together to drive the rhythm. His system is a little soft in the treble, and doesn’t have as much dynamic impact as he’d like.

Let's look at the sonic differences between two inexpensive and similarly priced DACs and see how each would—or wouldn't—fit into these systems.

DAC #1 has terrific bass: tight, deep, driving, and rhythmically exciting. Unfortunately, its treble is a little etched, grainy, and overprominent. DAC #2's best characteristics are a sweet, silky-smooth treble that has a complete lack of hardness, grain, and etch, and produces no fatigue in the listener. Its weaknesses, however, are soft bass and limited dynamic range. It doesn't have DAC #1's driving punch and dynamic impact with drums.

You can guess which DAC would be best for each system and listener. DAC #1 would only exacerbate the brightness Listener A's system already exhibits. Moreover, the additional grain would be more objectionable with violins and voices. DAC #2, however, would tend to soften the treble presentation in Listener A's system, providing much-needed relief from its relatively relentless treble. Moreover, the sonic qualities of DAC #1—dynamic impact and tight bass—are less musically important to Listener A.

Conversely, Listener B would be better off with DAC #1. Not only would DAC #1's better dynamics and tighter bass better serve the kind of music Listener B prefers, but his system could use a little more sparkle in the treble and punch in the bass.

Which DAC is "better"? Ask Listener A after she's auditioned both products in her system; she'll think DAC #2 is greatly superior, and wonder how anyone could like DAC #1. But Listener B would find her choice lacking rhythmic power, treble detail, and dynamic impact. To him, there's no comparison: DAC #1 is the better product.

Although I've exaggerated to make my point, this example shows how personal taste, musical preference, and system matching can greatly influence which digital products are best for you. The only way to make the right purchasing decision is to audition the products for yourself. Use product reviews in magazines to narrow your choice of what to audition, reading reviewers' descriptions of particular products to see if the type of sound described is what you're looking for. But don't buy a product solely on the basis of a product review—a reviewer's system and musical tastes may greatly differ from yours. You could be Listener A, and be reading a review written by someone with Listener B's system and tastes.

Use reviews not as sources of absolute truth, but as guides pointing you toward products you might want to audition yourself. You're going to spend many hours with your decision, so listen carefully before you buy—it's well worth

the investment in time. Moreover, the more products you evaluate and the more careful your listening, the sharper your listening skills will become.

It's important to realize that, generally, the less a DAC costs, the more pronounced its sonic signature. At the very highest levels of digital playback, the sonic trade-offs are much less acute—the best products have fewer shortcomings, making them ideal for all types of music—and two *perfect* DACs would sound identical. But, of course, there's no such thing as a perfect audio component.

Still, a significant factor in how good any DAC or CD player sounds is the designer's technical skill and musical sensitivity. Given the same parts, two designers of different talents will likely produce two very different-sounding products. Consequently, it's possible to find skillfully designed but inexpensive products that outperform more expensive products from less talented designers.

Higher-priced products are not necessarily better. Don't get stuck in a specific budget and audition products only within a narrow price range. If an inexpensive product has received a rave review from a writer you've grown to trust and the description of its sound matches your taste, audition it—you could save yourself a lot of money. If you decide not to buy the product, at least you've added to your listening database, and can compare your impressions with those of the reviewer.

In addition to determining which digital products let you enjoy music more, there are specific sonic attributes you should listen for that contribute to a good-sounding digital front-end. How high a priority you place on each of these characteristics is a matter of personal taste.

Here are the musical and sonic qualities I look for in digital playback.

The first quality I listen for in characterizing how a digital component sounds is its overall perspective. Is it laid-back, smooth, and unaggressive? Or is it forward, bright, and "in my face"? Does the product make me want to "lean into" the music and "open my ears" wider to hear the music's subtlety? Or do my ears tense up and try to shut out some of the sound? Am I relaxed or agitated?

A digital product's overall sonic perspective is a fundamental characteristic that defines that product's ability to provide long-term musical satisfaction. If you feel assaulted by the music, you'll tend to listen less often, and for shorter periods. If the product's fundamental musical perspective is flawed, it doesn't matter what else it does right.

Key words in product reviews that describe an easy-to-listen-to digital source include *ease*, *smoothness*, *laid-back*, *sweet*, and *polite*. Descriptors such as *bright*, *vivid*, *incisive*, *etched*, *forward*, *aggressive*, *analytical*, and *immediate* all point toward the opposite type of sound.

There is a fundamental conflict between these extremes of presentation. DACs that sound smooth, laid-back, and polite may not actively offend, but they often lack detail and resolution. An absence of aggressiveness is often achieved at the expense of obscuring low-level musical information. This missing musical information could be the inner detail in an instrument's timbre that makes the instrument sound more lifelike. It could be the sharp transient attack of percussion instruments; a slight rounding of the attack gives the impression of smoothness, but doesn't accurately convey the sound's dynamic structure. Consequently, very smooth-sounding digital products often have lower resolution than more forward ones.

The other extreme is the digital product that is "ruthlessly revealing" of the music's every detail. Rather than smoothing transients, these products hype them. In a side-by-side comparison, a ruthlessly revealing product will seem to present much more detail and musical information. It will sound more upbeat and exciting, and will appeal to some listeners. Such a sound, however, quickly becomes fatiguing. The listener feels a sense of relief when the music is turned down—or off. The worst thing a product can do is make you want to turn down the volume, or stop listening altogether.

This conflict between lacking and over-emphasizing detail can be resolved by buying a higher-quality (and, often, higher-priced) DAC. I've found a few models that can present all the music, yet are completely unaggressive and nonfatiguing. This rare virtue is musically important. The digital front-end must walk a fine line between resolving real musical information and sounding etched and analytical.

Digital reproduction also has a tendency to homogenize individual instruments within the soundstage. This tendency to blur distinctions between individual instruments occurs on two levels: the instruments' unique timbral signatures, and the instruments' specific locations on the soundstage.

On the first level, digital products can overlay music with a common, synthetic character that diffuses the unique textures of different instruments and buries the subtle tonal differences between them. The music sounds as if it is being played by one big instrument rather than many individual ones, with a sameness to instrumental textures that prevents their distinctive characteristics from being heard.

The second way in which digital playback can diffuse the separateness of individual instruments is by presenting images as flat "cardboard cutouts" pasted on top of each other. Each instrument is not surrounded by an envelope of air

and space, the soundstage is flat and congested, and you can't clearly hear where the aural image of one instrument ends and the next begins. Instead of separate and distinct objects—i.e., instruments and voices—hanging in three-dimensional space, you hear a synthetic continuum of sound. Good digital playback should present a collection of individual images in three-dimensional space, with the unique tonal colors of each instrument and voice intact, and a sense of space and air between one image and the rest. This seems to be easy for analog to accomplish, but quite difficult for digital. A recording with excellent portrayal of timbre and space will help you identify which digital products preserve these characteristics.

Another important quality in digital playback is transparency of soundstage. This is the perception that the space in which the music is being performed is crystal-clear, open, and transparent. (The opposite of this is thick, congested, and opaque.) Soundstage transparency is analogous to looking at a city skyline on a perfectly clear day: Just as smog or haze will reduce the buildings' immediacy, vibrance, and visible detail, so too will sound-stage opacity detract from the musical presentation. With naturally miked recordings of acoustic music, transparency allows you to hear way back into the hall.

I've focused on these aspects of sound for the evaluation of digital products because they are the most common flaws in digitally reproduced music. You should also listen for the other aspects—treble grain, rhythm, dynamics, etc.—described in [Chapter 3](#).

Beyond these specifics, a good question to ask yourself is, "How long can I listen without wanting to turn the music down—or off?" Conversely, the desire—or even compulsion—to play one CD or digital file after another is the sign of a good digital front-end. Some components just won't let you turn off your system; others make you want to do something else. This ability to musically engage the listener is the essence of high-end audio. When judging digital front-ends, that should be the highest criterion.

Now that we know what to listen for in digital sources, let's look at some specific product categories, starting with disc players and working our way up to DACs.

Disc Players

Despite the decline of the Compact Disc, sales of which peaked in 2000, many of us have amassed enormous libraries of CDs over the years. Unless we rip them to a music server, we still need a machine on which to play those CDs. That's why new CD players will be available for many years to come. Many excellent CD players are available for under \$1000, including some legitimate high-end machines that sell for as little as \$500. This is, however, the lower limit of true high-end CD players. Below this level you enter the realm of mass-market products designed for maximum features and minimum manufacturing cost, not for musical performance.

If you opt for a CD player, look for one that has a *coaxial* digital output on an RCA jack. This will let you use the CD player as a transport if you upgrade to a separate DAC in the future. When you add a DAC to a CD player, you're replacing the CD player's integral digital-to-analog conversion circuitry with the presumably better-sounding D/A conversion circuitry in the outboard DAC. Notice that I said a *coaxial* digital output on an RCA jack. Nearly all mass-market CD players use a much inferior optical connector called *TosLink*. If you drive an outboard DAC with TosLink, you won't be getting the sound you paid for. Insist on a CD player with a coaxial digital output. That way, you can be assured of a clear upgrade path in the future.

Super Audio Compact Disc (SACD)

Another option is a CD player that is compatible with the *Super Audio Compact Disc* (SACD) format. About 15% of the CD players now on the market also offer SACD playback. Such players are often a bit more expensive, but allow you to enjoy the large library of great-sounding SACD titles. You can think of a CD/SACD player as simply a CD player with expanded capabilities; in features and functionality, a combination player isn't that different from a CD-only machine.

SACD provides up to 74 minutes of high-resolution digital audio in two- and six-channel modes. SACD's sound quality can be superb, and significantly better than that of CD. The first SACD titles were from Sony Music and a few audiophile labels, but as of early 2015, nearly 10,000 SACD titles are available from 130 different labels—up from 1500 titles in 2004. The ultimate source for information on SACD titles is the website sa-cd.net.

SACD offers the possibility (but not the requirement) of a “hybrid” disc containing two layers on the same disc. The layer closer to the disc's surface contains conventional 44.1kHz/16-bit audio that can be read by any CD player; the lower layer holds the high-resolution or High Density (HD) information (see Fig.7-1). This key element of SACD provides backward compatibility with the world's more than 1.7 billion CD players. The advantages of backward compatibility are manifold: a single inventory for record companies and music retailers, and the ability to play the same disc on CD or SACD players. Moreover, consumers needn't make a conscious decision to stop buying music in one format (CD) and start buying into a new format. The transition from CD to SACD can be gradual and painless, particularly as SACDs don't cost much more than CDs. Indeed, some consumers buying a hybrid SACD may be unaware that the disc is not a conventional CD.




Fig.7-1 A Super Audio CD contains a high-density layer and a conventional CD layer on the same disc. (Courtesy Sony Corporation)

Fig.7-2 shows the SACD's physical structure. It is, in essence, a CD with an additional layer of information-carrying pits embedded in its playing surface. This second, higher-density layer is semitransparent to the 750-nanometer wavelength of CD laser light, but reflective to the shorter, 650nm wavelength of the SACD player's second laser. The CD player's laser focuses all the way through the disc; the SACD player's laser focuses halfway through the disc, to the HD layer. A switch on SACD players allows the user to switch between layers, and to compare the CD and high-res layers. The CD-quality layer is sometimes called the "Red Book" layer, after the color of the cover of the official Sony/Philips publication describing the specifications of the Compact Disc.




Fig.7-2 Hybrid SACD physical structure. (Courtesy Sony Corporation)

More information is contained in the second layer because the information-carrying pits are about half the size of the CD layer's pits; these are spaced closer together than on a CD. This high-density layer contains a two-channel mix and, at the artists' discretion, a six-channel version of the music, along with the potential for text, graphics, and limited video information. You can select between stereo and multichannel playback from the same disc, depending on your tastes and the number of channels in your system.

There's one potential drawback to buying hybrid SACD titles; although you can rip the music on the disc's CD layer to a music server, you can't transfer the information on the SACD layer to a computer. The digital outputs on SACD

players don't give you access to the high-resolution bitstream, only the CD-quality signal. If you want to hear the title in SACD sound quality, you'll need to play the disc. But as we'll see in the next chapter, that's not entirely true; it is possible to download a growing catalog of titles in Direct Stream Digital (DSD), the encoding format used for the high-resolution tracks of SACDs.

Universal Disc Players

A universal disc player is compatible with virtually all optical-disc formats: CD, SACD, DVD-Video, DVD-Audio, and Blu-ray Disc (and the many encoding variations on each of those discs). Unlike the CD/SACD player, which can be considered a CD player that offers SACD compatibility, the universal machine includes complex video-processing circuitry, and must include a fundamentally different clock frequency that is based on a 48kHz sampling rate, not CD's 44.1kHz sampling rate. (The DSD encoding system of SACD is based on an integer multiple of CD's 44.1kHz and thus doesn't require a different clock frequency.) The addition of video circuitry and clocks of different frequencies greatly increases the machine's complexity and potential for audible degradation of the signal. Moreover, the more formats a machine can play, the more royalties its manufacturer must pay. This limits the amount the designer can spend on internal parts that affect sound quality, potentially compromising performance. Consequently, universal players rarely offer the best sound quality—with any format—for a given price level. If most of your listening is to CD, make sure that the universal player has acceptable sound quality when playing CDs. It would be a mistake to compromise performance of your primary medium just to gain compatibility with DVD and BD. If you're on a budget, want the simplicity of a single digital source component in your rack, or have an audio system that also functions as a home theater, a universal player will suit your needs. For those with more ambitious systems, a disc- or computer-based digital front-end combined with a separate universal disc player may be more appropriate. If you want the best of both worlds—compatibility with any format and great sound—consider one of the aftermarket modifications of stock universal disc players. The modifiers replace the stock machine's power supply, DAC, and analog circuitry with high-quality designs and implementation.

Blu-ray Disc

The Blu-ray Disc (BD) format was developed as a high-definition replacement for DVD-Video. By using smaller pits and a shorter-wavelength playback beam generated by a blue laser, a BD can store much more information on a 120mm optical disc than can a DVD: 25GB on a single-sided, single-layer disc, compared with 4.7GB for a single-sided, single-layer DVD. Blu-ray's greater storage capacity provides several intriguing possibilities for high-resolution digital audio.

You might know that the Dolby Digital format used in DVD is limited to 448 kilobits per second (kbps) for 5.1-channel audio—about one-tenth the bit rate of CD-quality audio. By contrast, BD's maximum audio bit rate is a whopping 24 megabits per second (Mbps), or more than 53 times the audio bit rate possible on DVD. This greatly increased bit rate allows BD to deliver high-resolution multichannel digital audio with unprecedented sound quality, either in conjunction with video or as an audio-only format. On concert or music BD titles, the audio stream might be encoded in the Dolby TrueHD or DTS HD Master Audio formats, which provide high-res multichannel digital audio with bit-for-bit fidelity to the original high-res master. These two formats, which are described in detail in [Chapter 11](#) (“Audio for Home Theater”), use a lossless compression system that roughly halves the signal's bit rate for storage on the disc.

The Blu-ray specification also includes an audio-only disc (no video) that contains up to 7.1 channels of uncompressed high-resolution digital audio with sampling rates of up to 192kHz and word lengths of up to 24 bits. Only a few specialty record labels have introduced titles in this format; the major record companies have so far ignored the massive sonic potential of Blu-ray Disc. As of early 2015, about 300 BD music titles are available in a mix of stereo-audio only, multichannel-audio only, and video with high-resolution audio.

Bass Management in SACD and Universal Players

The term *bass management* refers to a system in which bass is filtered out of some channels and directed to a subwoofer, if the system includes one. Bass management is necessary in multichannel audio systems, and is found in multichannel SACD players, universal disc players, and multichannel home-theater controllers.

To understand why bass management is needed, consider a recording containing full-bandwidth information (low bass as well as midrange and treble) in every channel. A good example is a recording of popular music in which the bass guitar and kick drum are at least partially positioned in the center channel. Most multichannel speaker arrays employ large, floorstanding loudspeakers in the left and right positions, and smaller speakers in the center and rear positions. Low frequencies from the bass guitar and kick drum would overload the small center speaker, introducing distortion and/or damaging the speaker. Bass management selectively filters bass from the channels you specify according to the kind of speakers that are in your system. That low bass can be directed to a subwoofer, if your system includes one.

If you've set up a home-theater receiver or controller and selected Large or Small from the menu for each of your speakers, you've worked with bass management. The Small setting simply engages a high-pass (low-cut) filter on that channel, preventing low bass from reaching the speaker. A typical cutoff frequency is 80Hz.

Ideally, bass management would be performed in the multichannel home-theater controller or preamplifier the player drives, and the connection between them would be digital. But because of copy-protection issues (more on this later), SACD players are connected to a controller or multichannel preamplifier via the player's analog outputs. The player manufacturers can't assume that the multichannel preamplifier has bass management, so they include it in the player.

SACD players typically have a selection of common loudspeaker configurations from which to choose; you simply engage one from the player's front panel. For example, Setting #1 may be ideal for the typical system with large left and right speakers, small center and surround speakers, and a subwoofer.

The bass-management setup controls are sometimes accompanied by a variable delay for each channel. You set the amount of delay for each speaker according to how far you sit from that speaker, so that the sound from all

speakers reaches you simultaneously. For example, if you sit 12' from the left and right speakers but only 6' from the rear speakers, delaying the sound by six milliseconds (6ms) will result in the coincident (simultaneous) arrival of the front and rear sounds.

How to Choose a Digital Source—Features and Specs

The wide range of features and capabilities available in digital source components affects not just sound quality, but also how the products function in your system. A good example is a DAC with a variable output that allows you to drive a power amplifier directly, without the need for a preamplifier.

When shopping for a digital front-end, remember that technical performance is secondary to sound quality. Manufacturers will often tout their products on the basis of some technical innovation, or because it uses the latest parts. Although interesting, such descriptions don't tell you how the product *sounds*. Many technical factors influence a component's musical performance; which parts it uses is only one of these. Don't buy a product just because it has a particular DAC—many DACs with excellent parts pedigrees just don't sound musical. Listen to the product and decide for yourself if it sounds good. Just as you wouldn't consider buying an amplifier based on how little THD it has, so you shouldn't choose a digital component because it features parts used successfully in other, highly regarded products.

Another claim you should consider with suspicion is a processor's *jitter* specification. Jitter degrades sound quality by introducing timing variations in the critical digital-to-analog conversion process. Some manufacturers make unverifiable claims about their products' jitter performance, sometimes even picking numbers out of the air. Not only are there no standards for jitter specifications, jitter is very difficult to measure—a condition that has led to the current confusion in the marketplace. When shopping for digital components, forget the marketing hype—just listen.

Let's look at the features and specs you should consider when choosing digital source components.

Disc Player and DAC Features and Specs

In choosing a disc player, the first decision to make is which formats you intend to play, a topic covered earlier in this chapter: CD player, CD/SACD player, or universal disc player? After you've chosen, you'll need to consider which features and functions are important to you. Some of today's CD players sport a host of features and capabilities unimaginable to the format's creators.

A basic CD player will have one pair of analog outputs on RCA jacks and, usually, a coaxial digital output on an RCA jack and/or an optical digital output on a TosLink jack. The digital output allows you to upgrade your digital front end by turning the disc player into a transport that feeds an outboard DAC.

The outboard DAC converts a digital datastream from a transport, music server, or other source to an analog signal that is fed to your preamplifier. Outboard DACs range in price from \$150 to \$35,000, but many good-sounding models can be had for under \$1000. The most basic DAC has one USB digital input and one pair of unbalanced analog outputs. More complex DACs have multiple digital inputs, digital outputs, balanced analog outputs, polarity-inversion switching, and sometimes a volume control.

Fig.7-3 shows examples of DACs of varying complexity. The example at the top is about as simple as a DAC can get; it plugs into your computer's USB port, replacing the computer's internal DAC with a higher-quality outboard DAC. The analog output is on a 3.5mm stereo mini-jack that can drive headphones or, with a mini-plug-to-RCA cable, connect to a home stereo system. The DAC on the bottom offers multiple digital inputs, variable output level, user-selectable digital filters, a *polarity-inversion switch*, and remote control.



Fig.7-3 DACs range from simple plug-in devices to elaborate multi-input machines. (Courtesy Audioengine [top] and Berkeley Audio Design [bottom])

I've combined disc player and DAC features and specifications because the two component categories share so many similarities. Indeed, a disc player is simply a DAC with an integral disc transport.

Let's take a closer look at the features and options available with disc players and DACs.

Number and Types of Inputs: Many DACs offer a range of input connections, and virtually all DACs made since 2011 offer a *USB input* for connecting to a computer. Be sure that the DAC's USB input offers *asynchronous* operation, which means that the DAC runs off its own clock, not the computer's low-precision clock. The DAC's primary USB input for connecting a computer is on a Type B connector, usually located on the rear panel. A few DACs also offer a front-panel Type A USB jack that allows you to play music directly from a USB memory stick or hard-disk drive. These DACs sometimes include the ability to decode different file formats, such as *FLAC*, *AIFF*, *MP3*, etc., and can play music files directly from a storage medium without the need for a computer. (File formats are explained in the following chapter.) These different USB types are being replaced by the new Type C connector, which is reversible both in direction and up/down orientation, and makes a click when fully inserted in a jack.

Most DACs provide *S/PDIF* input on an RCA jack, or sometimes on a *BNC jack*. A BNC jack is technically superior in this application, but very few sources provide BNC output (Fig.7-4 shows a range of digital inputs, including two BNC jacks in the middle of the rear panel.) The next most common input type is

AES/EBU, recognizable from the three-pin XLR jack. AES/EBU is a professional version of S/PDIF that's carried on a balanced cable. Virtually all DACs have an optical input on a TosLink jack for compatibility, but TosLink is sonically inferior to the other input types. A very few DACs have ST-Type optical connection, an interface popular in the 1990s. As mentioned earlier, DACs with networking capability have an Ethernet jack for connecting to the network. (Networked audio is described in the next chapter.) When choosing a DAC, be sure that it has the type of inputs that will accommodate the output from your digital source components.



Fig.7-4 A DAC's rear panel showing balanced and unbalanced outputs, along with a variety of input types. (Courtesy Berkeley Audio Design)

High-Resolution PCM Decoding: Virtually all of today's DACs, and the DAC sections of disc players that offer digital inputs, can decode digital bitstreams with resolution beyond CD's 44.1kHz sampling frequency and 16-bit word length. Some DACs are limited to 96kHz/24-bit; others can decode any sampling frequency up to 192kHz and word lengths up to 24 bits. The first generation of USB DACs was limited to 96kHz because of the available chipsets; later models can accommodate up to 192kHz, the highest frequency of any currently available source material.

DSD Decoding: Driven by the increasing availability for download of files in Direct Stream Digital, the encoding format of SACD, DAC manufacturers have recently begun to include the ability to decode DSD. Note that downloaded DSD files can be standard-rate DSD (2.8224 megabits per second, the same as SACD), *double-DSD* (5.6448Mbps), *triple-DSD* (8.4688Mbps), and even *quad-DSD* (11.2928Mbps). You'll also see standard-rate DSD referred to as DSD64, because the sampling frequency is 64 times that of CD; double DSD is DSD128; and so forth. DACs with DSD decoding vary in their ability to accept and decode these higher-speed DSD files.

Note that virtually all DSD-compatible DACs convert the DSD bitstream into a PCM signal for digital-to-analog conversion. As explained in detail later in this chapter, one of DSD's advantages is that it can be converted to analog by a very simple DAC, with no need for the complex and sonically degrading "decimation filters" required for PCM conversion. DACs that convert DSD signals to PCM obviate one of DSD's advantages.

DAC manufacturers can add DSD compatibility in one of four ways. The first is simply to use the DSD-to-PCM conversion circuit built into most DAC chips. This approach costs the manufacturer nothing, but introduces sonic compromises because the DSD-to-PCM conversion is performed on the same chip as the critical digital-to-analog conversion. The second way of making a DAC DSD-compatible is to add some circuitry to perform the DSD-to-PCM conversion outside the DAC chip. The third method is to convert DSD files to PCM in a computer's music-management software on the fly during playback. This step bypasses the DSD-to-PCM conversion in the DAC. An advantage to this approach is that the DSD-to-PCM conversion algorithms can be upgraded with an update of the music-management software, an option not possible with the DSD-to-PCM conversion built into the DAC's hardware. Note that you can also convert DSD to PCM in the computer and store the file as PCM data, making it unnecessary for the music-management software to perform the conversion on the fly. The fourth method of making a DAC DSD-compatible is to build two independent decoding paths within the case, one for PCM and one for DSD. This architecture keeps the DSD datastream in the DSD domain, and is sometimes called native DSD conversion.

A more technical explanation of how DSD works appears later in this chapter.

Volume Control and Source Switching: For those who listen only to digital sources (no analog tuner or turntable), the disc player or DAC with a volume control is ideal. This feature allows you to drive a power amplifier directly, with no need for a traditional analog preamplifier. This saves money—on both the preamp and an additional pair of interconnects—and has the potential to sound better because fewer electronic components are in the signal path.

A related feature is called *source switching*. A disc player or DAC with source switching has multiple digital inputs, allowing you to listen to different digital sources through the DAC or disc player's DAC section. A few products even include one or two analog inputs in addition to digital ones, completely doing away with the need for a traditional preamplifier in most systems. It is unlikely,

however, that the preamplifier section built into a DAC can match the sonic performance of a standalone preamplifier.

The volume control in disc players and DACs can be implemented in the analog or the digital domain. That is, the analog audio signal can be put through a standard volume-control knob, as is found on a preamplifier, or the volume can be adjusted by performing mathematical operations on the digital data representing the music. Before deciding on a disc player or DAC with volume control, you should know the trade-offs inherent in each approach.

An analog volume control is a mechanical device that subjects the audio signal to a varying electrical resistance. That resistance, and the contacts between the resistance and parts inside the volume control, can slightly degrade the signal—no volume control is perfectly transparent. Moreover, the volume control can introduce small channel-balance errors at certain volume settings. For example, when the volume is turned very low, the left channel may be half a dB louder than the right. This situation could reverse as the volume is turned up, due to manufacturing tolerances in the volume control. Better volume controls use a switched-resistor network (described in [Chapter 4](#), “Preamplifiers”), but DACs with this feature are rare and expensive.

A digital volume control doesn't suffer from these problems but has its own drawbacks. Each 6dB of attenuation (reduction in volume) from the maximum setting throws away one bit of resolution. A low volume setting (say, 30dB of attenuation) is equivalent to discarding five bits. If you had true 20-bit resolution in your D/A converter, you'd be listening to 15-bit audio instead of 20-bit. The lower the volume setting, the greater the loss in resolution. This shortcoming of digital volume controls is much less significant in today's products than in those of ten years ago. As computing horsepower has increased, designers have been able to implement digital volume controls with greater precision. Implemented correctly, a digital volume control can be sonically benign at most listening levels. My current reference DAC has digital volume control, and I hear no loss of resolution at a typical playback volume.

A third type of volume control is a hybrid of analog and digital attenuation. As you turn down the volume from full scale, the attenuation is performed in the digital domain up to 5.9dB of attenuation. At the next step down, 6dB of attenuation, the digital signal is restored to full scale and 6dB of analog attenuation is engaged, via relays and a switched-resistor network. At, say, 9dB of attenuation, 6dB of the attenuation will be digital and 3dB analog. This innovative method provides a wide range of attenuation with practically no loss in resolution.

One final note about digital-domain attenuation: DACs with this feature often sound very slightly better with a few dB of attenuation compared with full-scale output. The reason is that the digital filter has an easier time processing signals slightly below full scale.

Remote Control: DACs with a volume control should be supplied with a remote control so that you can adjust the output level without leaving the listening seat. This feature is more important if your DAC drives a power amplifier directly rather than through a preamplifier.

Selectable Digital Filters: Some DACs provide multiple digital filters, selectable by a front-panel button. These filters all have different performance trade-offs, which results in each filter offering a slightly different sound. The DAC's default filter is usually the most technically correct, with the other filters representing a departure from "textbook" performance but offered as a way of fine-tuning the sound. For example, in one DAC with six filters, the first filter results in the best measured performance but the brightest sound; each successive filter sounds slightly softer and more forgiving in the treble.

Clock Input: More sophisticated and expensive disc players and DACs have an input that allows them to be clocked by an external box that contains a precision clock of presumably higher precision. This feature is found only in more expensive products.

HDCD Decoding: Some disc players and DACs can decode CDs recorded using the *High Definition Compatible Digital (HDCD)* process. HDCD is explained later in this chapter.

Balanced Outputs: Many disc players and DACs have *balanced outputs* on XLR jacks, in addition to unbalanced outputs on RCA jacks. If you have a fully balanced preamplifier, consider a DAC with balanced output jacks.

As explained in [Chapter 10](#), a balanced signal is carried on three conductors: a ground wire and two signal-carrying conductors. The two signals are identical but opposite in polarity; i.e., they are mirror images of each other. When one is swinging positive, the other is swinging negative. However, at the end of this chapter, under "How a DAC Works," you'll see that not all DACs with XLR jacks are truly balanced.

Digital Output: This feature allows you to send the digital signal to a digital recording device. A digital output is an important feature if you own a digital recorder—without it, you’ll need to unplug the digital signal going into your DAC and run it to the digital recorder every time you want to record through the recorder’s digital input.

Polarity-Inversion Switch: A switch on the front-panel or remote inverts the absolute polarity of the signal to match its original polarity. To find correct absolute polarity for a particular recording, simply throw the switch back and forth to find the position that sounds better (if there is an audible difference). A polarity-inversion switch is often marked “180°,” to indicate polarity reversal. Sometimes called (incorrectly) Phase.

Internet Radio: A relatively few DACs include the ability to receive and decode *Internet Radio* stations. Internet Radio, also called webcasting, allows you to hear nearly any radio station in the world through an Internet connection. Just as with traditional terrestrial radio, Internet radio is a streaming service; you can’t pause or replay. Internet radio is highly compressed and of low audio fidelity; a typical bitrate is 128kbps (contrasted with 1.41Mbps for CD).

Networking Capability: Some DACs have built-in networking capability that allows them to connect to a local area network (LAN) and access music files from other devices on the network, such as a *network-attached storage* (NAS) drive. This network can be wired or wireless, with the wired connection on an Ethernet jack. An app supplied by the DAC manufacturer allows you to manage your music library and select music to play. Some DACs with networking capability require that you purchase an additional component to fully realize the networking and music-storage features. While such DACs are sometimes called “renderers,” a renderer can also convert audio streamed from a network to a universal format such as S/PDIF.

Digital Interfaces

At first glance, it would seem that the method of getting the digital datastream from a source to a DAC would have no effect on sound quality, provided that the interface doesn't introduce errors. In other words, if the ones and zeros are the same at the receiving end as at the transmitting end, the sound should be the same. It's a trivial matter to prove that digital interfaces don't introduce bit errors, but they nonetheless exhibit an analog-like variability in sound quality. Different interface types sound different from one another, and different brands of digital interconnects all exhibit their own sonic flavors. To some skeptics, reports that digital cables vary in sound quality is the height of audiophile lunacy. According to the conventional wisdom, if the ones and zeros are the same, the sound must be the same. After all, a computer program runs identically, whether from a computer's internal hard drive or from an external drive connected by a USB cable. As described later in this chapter ("Jitter in the Digital Interface"), the analog-like variability among digital interfaces and cables is the result of jitter, or timing imprecision.

Let's look at the range of digital interfaces seen in today's products.

The S/PDIF Digital Interface

Ever since the earliest CD players, the *Sony/Philips Digital Interface Format (S/PDIF)* has been the conventional interface found on all types of digital products: the ubiquitous “digital out” RCA or optical jack found on components from disc players to personal computers. As described earlier, transports and DACs come fitted with a variety of connection methods: AES/EBU, coaxial, TosLink optical, or BNC. All of these interfaces transmit virtually the same S/PDIF format, but in different ways.

Digital interfaces can be divided into two categories: electrical and optical. In an electrical interface, electrons carry the signal down copper or silver wire. An optical interface transmits light down a plastic or glass tube.

Fig.7-5 shows the three common interconnect types: coaxial, AES/EBU, and TosLink. Coaxial and AES/EBU are both electrical; TosLink is optical. The most common type of connection, coaxial, is carried on an RCA cable. This is the electrical connection found on virtually all CD transports, most good CD players, and other consumer digital audio products such as digital preamplifiers, room-correction systems, and digital recorders. A variation on the coaxial connection is the BNC cable, which is used in only a few products. Though BNC is mechanically, electrically, and sonically superior to RCA, it never caught on. Still, some manufacturers offer BNC connections. If your components offer BNC connection, use it.

TosLink is the low-cost optical interface promoted by mass-market audio manufacturers as an alternative to coaxial connection. A trademarked name of the Toshiba Corporation, TosLink is more properly called EIAJ Optical, after the Electronics Industries Association of Japan. The major electronics companies had two good reasons for trying to convert the world to TosLink. First, TosLink jacks and cables are cheaper to make than coaxial jacks and cables. Second, TosLink connection makes it easier for components to meet Federal Communications Commission (FCC) requirements for radiated noise. An electrical signal traveling down copper wire (such as the S/PDIF signal in a coax connection) throws off radio-frequency (RF) noise that could interfere with radio and television transmission. The FCC will simply ban products that don't meet its criterion for radiated noise. Because TosLink sends the signal as light down glass or plastic fiber, it produces no radiated noise.

The ST-Type optical interface, developed by AT&T for telecommunications, was popular in the mid-1990s but now rarely appears on new audio products. ST-

Type connection transmits the optical signal down wide-bandwidth glass fiber instead of TosLink's plastic light path (some higher-quality TosLink cables also use glass). ST's locking bayonet connector ensures a good junction between the cable and optical transmitter and receiver.



Fig.7-5 Comparison of different digital interface terminations. From left to right: coaxial S/PDIF on RCA jack, AES/EBU on XLR jack, TosLink optical. (Courtesy AudioQuest)

Many products offer the AES/EBU interface, which is carried on a balanced line terminated with three-pin XLR connectors. Of the three conductors in a balanced signal, one is ground, one is the digital signal, and the third is the digital signal inverted. AES/EBU benefits from all the advantages of balanced lines (described in [Chapter 10](#)) and is transmitted at 5V, compared to S/PDIF's 0.5V. AES/EBU is usually the preferred interface because of this higher voltage, which confers greater timing precision, as well as greater noise rejection by virtue of the balanced connection.

Virtually all transports and processors have RCA jacks, and most include TosLink for compatibility. AES/EBU is usually found only on more expensive components, or is offered as an option on midpriced digital equipment. TosLink is by far the worst interface, mechanically (the physical connection between cable

and jack), electrically (it has the lowest bandwidth), and thus sonically. TosLink tends to blur the separation between individual instrumental images, adds a layer of grunge to instrumental textures, softens the bass, and doesn't provide the same impression of "black" silence between notes. Better results can be achieved with a high-end glass-fiber TosLink cable, but my advice is to forget about TosLink unless you have no other choice.

Some digital front-ends will benefit from an optical interface because there is no ground connection between the source and DAC, as exists with an electrical cable. Sonically degrading high-frequency noise on a digital source's ground can contaminate the DAC's ground through an electrical connection. Because an optical interface transmits the signal without an electrical connection, there is no chance of coupling ground noise between the two components.

In theory, the electrical interfaces will work best because they have the widest bandwidth. The transmission of digital data from a source to a DAC should be over the widest possible bandwidth link, for low jitter. TosLink has a bandwidth of 6MHz (6,000,000Hz), and electrical interfaces have a potential 500MHz bandwidth (if implemented correctly).

The best way to choose an interface is to listen to the various types. When comparing electrical and optical interfaces, be sure to disconnect the electrical interface from the DAC when listening to the optical interface. The advantage of keeping the source's ground isolated from the DAC ground is lost if the electrical cable is still connected between the source and DAC, even if it's not active.

A primary reason for poor-sounding sources and interfaces in electrical connections is a mismatch between the source's output impedance, the cable's characteristic impedance, and/or the DAC's input impedance. Such an impedance mismatch can cause reflections of the signal in the cable, which introduces jitter in the datastream. The impedance specification for S/PDIF is 75 ohms ($\pm 5\%$), and for AES/EBU, 110 ohms ($\pm 20\%$). Manufacturers of digital products should strictly adhere to these impedance standards, although many do not.

Finally, it should be noted that recent improvements in DACs have largely reduced the sonic differences between interface types, and between specific models of interface cables. Today's DACs are much better at rejecting jitter, and some even employ technologies that completely isolate the DAC from interface jitter. Keep this in mind: The greater the differences you hear between interfaces, the worse the DAC is at rejecting interface-induced jitter.

The USB Interface

The Universal Serial Bus (USB) was originally designed to connect computer peripherals to a PC, replacing the serial and parallel ports on earlier-generation machines. The popularity of PC-based music servers has prompted the high-end audio industry to adapt the interface for digital audio, specifically to connect the output of a computer-based music server to a DAC. Virtually all modern DACs offer USB compatibility.

Connecting a Mac computer to a USB DAC is usually straightforward: You simply connect the USB cable, and the computer (the “host,” in USB parlance) recognizes the DAC (the “device”). Windows-based computers often require that you download and install the USB driver for the particular DAC.

Most DACs with USB inputs available as of this writing conform to the USB 2.0 protocol. Although USB 2.0 has a bandwidth of 480Mbps (far faster than any high-resolution digital audio signal), some older DACs (and a few modern ones) that adhere to the USB 1.0 standard are limited to 96kHz sampling frequency. That’s not a limitation of USB itself, but rather of the chipsets DAC manufacturers choose for their lower-cost products. One advantage of USB 1.0 DACs is that they don’t require a driver download when connected to Windows-based computers.

A USB cable for audio has a Type A connector on the computer end and a Type B connector on the DAC side (Fig.7-6). The new USB 3.1 specification replaces both of these with small Type C connectors, making the cable reversible. Because the Type A and B connectors are so common (they’re now nearly 20 years old), it will be quite some time before they’re phased out entirely.

In virtually all digital interfaces, including S/PDIF and AES/EBU, the source component (e.g., the CD transport or computer) is the master clock to which the receiving device (the DAC) must lock. This method is problematic, for two reasons. First, the clock “recovered” by the DAC from the datastream is generally high in timing variations, aka jitter, that degrade sound quality. Second, the clock generated by a computer is far from optimal; it will be noisy and high in jitter. Early USB DACs operated in this way, called *adaptive mode*, which severely degrades sound quality. The industry quickly responded to create *asynchronous mode* USB, in which the DAC’s critical clock runs independently of the PC’s clock. This technique greatly reduces jitter and results in better sound. Asynchronous mode transforms the USB interface into a high-quality interface, provided that the USB interface is implemented correctly. More than any other

interface, USB has the widest variability in sound quality. Some otherwise excellent DACs are compromised by a poorly designed USB interface.



Fig.7-6 A USB cable for audio has a Type A connector on the end that connects to a computer (lower image) and a Type B connector on the end that connects to the DAC (upper image). (Courtesy AudioQuest)

If you own an older DAC that doesn't offer a USB interface yet you want to drive the DAC with a computer's USB output, you can buy a USB-to-S/PDIF converter. These devices are small and inexpensive, but generally don't offer high-end sound. Some USB-to-S/PDIF converters do deliver state-of-the-art performance, but they're much more expensive than the many "convenience" devices on the market. These better converters can actually exceed the sonic performance of any USB interface built into a DAC because they isolate the sensitive DAC circuits from the computer, as well as reclock the signal with high precision.

Note that because inexpensive USB DACs get their power from the computer via the USB bus, they don't need to be plugged into a 120V AC wall outlet. Others have either an integral power supply, or an external "wall wart" supply that sends DC power to the DAC.

Finally, here's a hint regarding USB cable length. Unlike speaker cables and analog interconnects, where shorter is always better, the best-sounding USB cables are 1.5 meters long.

Ethernet

The *Ethernet* interface (Fig.7-7) is used in networked audio systems to carry digital audio, along with metadata and control signals that enable all the devices on the network to communicate with each other and share files. The following chapter describes networked audio systems.




Fig.7-7 An Ethernet cable in a networked audio system can carry audio as well as metadata and control signals. (Courtesy AudioQuest)

The I²S Interface

The *I²S interface* (pronounced *I-squared-S*) was originally developed by Philips as a means of transporting audio data and clocks between chips inside a CD player. A few high-end manufacturers have tapped into this bus to link separate components to each other. I²S has the advantage of carrying the clock on a separate line rather than combining it with the audio data, as with S/PDIF.

The I²S interface is usually implemented on a so-called 13W3 cable, originally designed for the computer industry for transmitting component-video and control signals. It's also possible to carry I²S signals on other cable types, including RJ-45 and HDMI (High-definition Digital Multimedia Interface).

There was a movement in the early 1990s by high-end manufacturers to make I²S a standard, but only a few companies offered the interface. I²S then fell so far out of favor that I removed references to it in the third edition of this book. It's been interesting to watch its return to the market.

DSD over PCM (DoP)

The increasing popularity of the DSD format for high-resolution downloads presented a challenge to hardware manufacturers and music-management software developers: how to transmit DSD over a digital audio interface that had been designed to carry data encoded as a Pulse Code Modulation (PCM) signal. Although each manufacturer could have designed a custom solution specific to a particular operating system or computer platform, that approach would have not only added greatly to the products' costs, but also precluded compatibility between products from different manufacturers. The industry needed a universal standard so that consumers could buy products knowing that those products would work with many others on the market.

The solution was to format DSD data so that it looks like PCM audio to the computer's audio-interface driver. The interface is thus called *DSD over PCM*, or DoP. Disguised as PCM data, the DSD data are transmitted over the USB interface, then reconverted to DSD data by a DSD-compatible DAC. Note that the music-management software and the DAC must each support DoP.

DoP was developed by Channel D Software and dCS, with the first public demonstration in June 2011, by Channel D Software. Later that year DoP's developers offered the technology royalty free to the industry. Many hardware manufacturers and music-management software developers have adopted the standard, which continues to be refined, by consensus, of the supporting manufacturers.

In addition to making DSD playback more affordable and universally available, DoP also protects your loudspeakers from damage. Without DoP, a raw DSD datastream fed to a PCM DAC by mistake would cause the DAC to output speaker-damaging, full-scale white noise. But a DoP-formatted signal sent to the DAC would result in a very low noise level (-48dB).

Proprietary Digital Interfaces

Some digital sources employ a proprietary high-resolution digital output that allows the company's SACD player or other high-res source to connect digitally with the company's other products. An example is the Meridian High-Resolution (MHR) Smart Link, which features proprietary encryption to prevent illegal copying of the datastream. Software in Meridian's controllers encrypt the data for output on the *MHR Smart Link* jack, which can be received and decoded by Meridian's other products, notably its digital loudspeakers. The MHR interface also carries information about the disc content, such as whether the disc is stereo or multichannel, and the audio encoding format. This information allows the MHR-equipped controller to automatically invoke the correct decoding and DSP modes.

Some Sony SACD players offer high-resolution, multichannel digital output on an optical *i.LINK* connector, which can connect to some of Sony's other products as well as to outboard high-end DACs featuring i.LINK input. I.LINK is also Sony's trade name for the IEEE1394 interface, aka FireWire. This high-res digital signal is copy-protected by the Digital Transmission Content Protection (DTCP) system.

Another example of a proprietary digital interface is PS Audio's use of the I²S format, carried on an *HDMI* interface. Note that all of these interfaces are nonstandard and specific to the companies that introduced them. Consequently, components that include these interfaces won't connect via the proprietary connections to products made by other manufacturers.

Wireless Connection

So far, I've talked about getting digital audio from one place to another via wires, whether electrical or optical. But increasingly, manufacturers are providing wireless connectivity in their digital components. There are no universal standards for wireless connection, but a few have emerged as leaders.

The most common wireless interface is *Bluetooth*, commonly found in mass-market products such as desktop audio systems and audio/video receivers and primarily used to stream audio from a mobile device to the playback system, and for whole-house distributed audio. Bluetooth is simpler than WiFi because it transmits signals directly between components rather than going through a network and router. Unfortunately, Bluetooth's sound quality is compromised because the limited bandwidth requires lossy data compression that sacrifices sound quality for a lower bitrate. Two initialisms associated with Bluetooth are *A2DP* and *aptX*. Advanced Audio Distribution Profile (A2DP) is a one-way communication format that can employ a number of lossy codecs, including Apple's popular AAC. Most Bluetooth devices support this protocol. Some also support aptX, which offers a more sophisticated codec for better sound. Promoters of aptX suggest that it offers "near CD-quality" wireless streaming. Bluetooth devices are classed according to the distance over which they will reliably work. Class 1 has a theoretical range of 100 meters, but the actual distance within which a Class 1 device can be used is a fraction of that. Class 2 devices are specified at 30 meters, but again, that's an optimistic figure.

An alternative wireless connection is Apple's *AirPlay*, a system familiar to iOS users. Unlike Bluetooth, AirPlay is based on WiFi and consequently requires that any devices that share content be connected to a wireless network. AirPlay will stream music stored on a Mac or mobile device, or via services such as Tidal, to any AirPlay-compatible receiver. AirPlay uses Apple Lossless compression, which delivers bit-for-bit accuracy between the source and receiver—a big advantage over Bluetooth. Note, however, that AirPlay won't carry high-res audio; it downconverts higher sampling frequencies to 44.1kHz. Finally, AirPlay is not an open standard, but proprietary to Apple. If you use an iPhone to stream audio to your playback system and decide to switch to an Android phone, all the AirPlay hardware you've invested in won't communicate with the new phone.

Some DACs offer WiFi streaming of uncompressed, high-resolution audio via a proprietary app. With these DACs, wireless connection can confer a significant sonic advantage over an electrical connection such as USB. When a DAC and

computer are connected by an electrical cable, the computer's very noisy ground is connected directly to the DAC's ground, potentially degrading the sound quality. With no physical connection between computer and DAC, the DAC is isolated from the computer's noise. When I've compared high-resolution, uncompressed WiFi streaming to a USB connection, WiFi has sounded smoother, cleaner, and more transparent than USB. Note, however, that the quality of connection between computer and DAC is subject to the robustness of the WiFi network. A network with low signal strength could result in dropouts in the music.

I think of most wireless connection as a "lifestyle" technology; it's convenient, it lets you easily connect devices and distribute music to other rooms in your home, but it rarely offers the ultimate in sound quality unless it's done right.

Jitter in the Digital Interface

Jitter refers to sound-degrading timing inaccuracies created in the process of converting a digital to an analog signal. A clock in the digital processor controls when the digital samples that represent the music are to be converted to analog. If this clock isn't perfectly precise, sound quality suffers.

Think of the bass drum in a marching band as the clock in a DAC, and of the band members as the digital samples representing the music. Each row of musicians is one audio sample. On each drum beat, the musicians take one step forward in unison. Now consider what would happen if the drummer got sloppy, leaving too much time between some beats and not enough between others. Without a precise timing reference, the distances between the rows of musicians would no longer be uniform.

In digital sound reproduction, each clock pulse triggers the conversion of a digital sample to analog form. If these clock pulses aren't perfectly uniform, we hear the resulting timing imprecision—the very definition of jitter—as a degradation of sound quality. Jitter causes the samples representing the music to be put back together with slightly staggered spacing, creating distortion in the audio waveform. Jitter is audible as reductions in the sense of space (the recorded acoustic sounds smaller and drier), the impression of three-dimensional objects (musical instruments) in that space, the liquidity instrumental textures, and the firmness of the bass. Just as the marching-band musicians can't step together if the drummer isn't perfectly precise, digital audio samples can't be correctly reconstructed into music if the clock is jittered. To give you an idea of how precise the clock must be, timing errors of as little as 10 picoseconds (0.00000000001 second)—about the time it takes light to travel 1/10"—are audible. (A more technical explanation of clock jitter appears in [Appendix C](#).)

The S/PDIF interface is fundamentally flawed, in that the clock is carried within the audio data (though it's more technically accurate to say that the audio data are embedded in the clock). This requires that the receiving device lock to that clock with a phase-locked loop (PLL) and generate a new clock based on the incoming clock. PLLs are imperfect devices; they tend to track and pass along rather than reject variations in the incoming signal. The clock "recovered" by the PLL becomes the timing reference for the critical digital-to-analog conversion stage. This is the point where jitter matters—in the clock that controls when each digital sample is converted to an analog value in the DAC chip. If that clock isn't perfectly precise, the digital samples that represent the analog signal's amplitude

are reconstructed with slightly staggered spacing between samples, resulting in tiny amplitude errors in the analog waveform. Such micro-errors in amplitude don't exist in nature, which perhaps explains why our brains are highly sensitive to the phenomenon, and why transports and interfaces affect sound quality.

This interface-induced jitter isn't random in nature; it's directly correlated with the audio signal being transmitted. In other words, the clock is modulated by the music carried by the interface. The recovered clock is jittered at exactly the same frequency as the audio signal traveling down the interface. Using a jitter analyzer, it's possible to listen to only the jitter component of the S/PDIF signal. Hearing *only* the jitter, you can still readily identify the piece of music being transmitted. This audio-correlated jitter ends up at the DAC's word clock, where it degrades sound quality.

Because a CD player has no need for an S/PDIF interface, it can theoretically have lower jitter than a separate digital source and DAC. One of the main sources of jitter is the digital interface that connects a CD transport to a DAC. Because the interface transmits the digital audio data and the clock along the same cable, jitter can be introduced there.

Outboard Clocks

In the early 2000s, several manufacturers introduced a new breed of product designed to reduce jitter and its harmful effects, and to be inserted in the digital signal path between a digital transport (or any other digital source) and a DAC. In theory, the jitter-reduction device accepts a jittered S/PDIF digital signal and outputs a less jittered digital signal.

These devices exhibited various degrees of efficacy. One fundamental problem was that the dejittered signal was sent to the DAC via an S/PDIF interface—which promptly reintroduced jitter into the signal.

Jitter-reduction devices fell out of favor after a few years, but manufacturers never stopped looking for solutions to the problem. The modern approach is to use a high-precision clock in a separate housing. When used with a transport and DAC, the clock's outputs supply both the transport and DAC with a precise clock (Fig.7-8)—the critical digital-to-analog conversion stage is no longer controlled by a clock recovered from the S/PDIF interface. These outboard clocks can be *extremely* precise. When used with a CD player, the clock's output becomes the master clock for the entire player, with the transport mechanism locked to the high-precision clock.

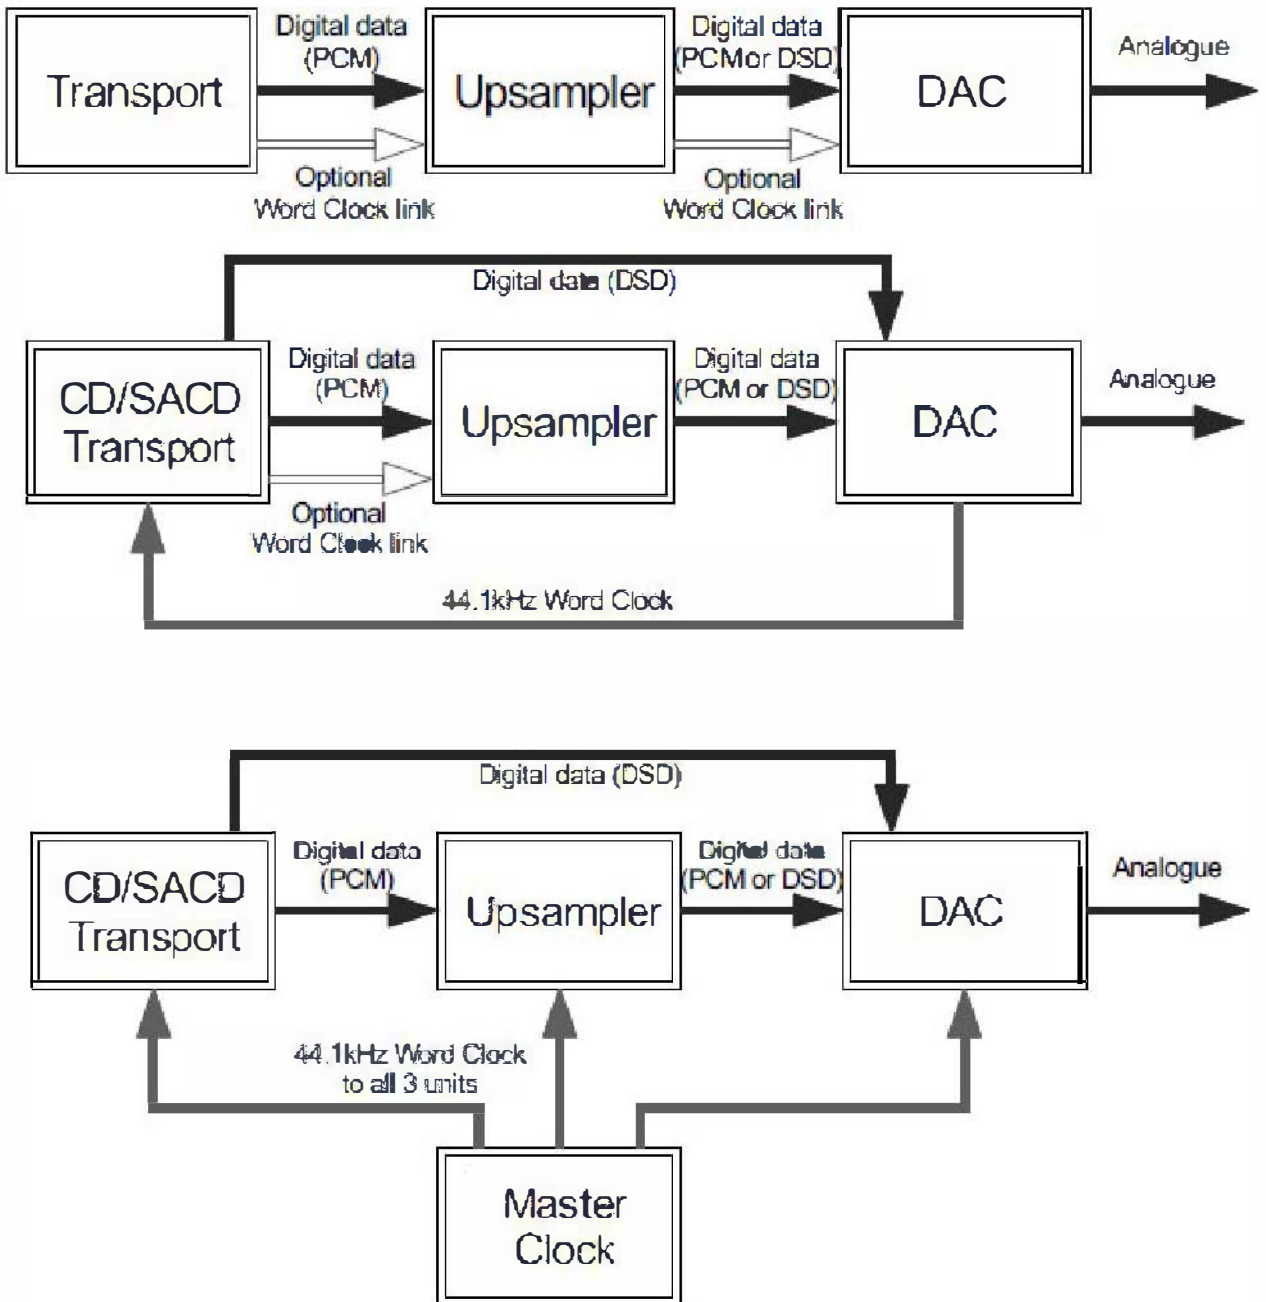


Fig.7-8 An outboard clock sends a precision timing signal to a transport and a DAC. This diagram shows three different clocking schemes: a transport, a separate upsampler, and a DAC. (Courtesy dCS)

It's a fascinating exercise to listen to the sonic effects of an outboard clock by turning it on and off. With the clock turned off, the CD player or DAC reverts to its own internal clock. It is thus possible to hear the effects of jitter with a simple button push. Lowering jitter sharpens image focus within the soundstage; expands the sound's depth, width, and layering; reduces the hardness on leading-edge transients such as upper-register piano notes; and, surprisingly, provides tighter and better-defined bass.

Perhaps the most ambitious clock is the Esoteric G-0Rb, a \$16,000 device that employs at its heart the element rubidium, which forces an oscillator to resonate at a precise frequency to create the timing reference. The G-0Rb is, in essence, an atomic clock designed solely to reduce jitter in digital audio reproduction. This might seem like overkill, but listening to an Esoteric transport and DAC with and without the G-0Rb was revelatory. The G-0Rb had a profound influence on the sound quality, both in specific performance attributes as well as in overall musical involvement. The sound with the clock had a much greater ease and naturalness, sounding less mechanical and artificial.

Another school of thought holds that an outboard clock is not the ideal solution because, before it reaches the DAC chip, the high-precision clock generated within the device must be output through electronics, sent down a cable, and received and processed by additional electronics, any and all of which might contaminate the clock. Designers who eschew outboard clocks instead work to create an extremely precise and noise-free clock right next to the DAC chip, to ensure that the clock remains uncorrupted at the DAC chip—the point where jitter matters.

Asynchronous Sample-Rate Conversion

Asynchronous sample-rate conversion is not to be confused with asynchronous-mode USB connection; the two are completely unrelated. Asynchronous sample-rate conversion refers to a technique in which digital audio data are converted to a different sample rate. The term *asynchronous* means that the input and output frequencies can run independently of each other and need not be related or locked together; the input frequency can be varied and the output frequency will remain the same. First used in professional applications in which multiple sampling rates were common (44.1 and 48kHz, for example), asynchronous sample-rate conversion has found its way into consumer audio products.

Asynchronous sample-rate conversion has the ability to remove jitter from the incoming signal. Theoretically, different transports, sources, and interfaces, each with an asynchronous sample-rate converter chip in its signal path, should be sonically identical (the chip is usually located in the DAC). The sample-rate converter chip is essentially an oversampling digital filter that interpolates new samples. This means that the output data are not identical to the input data, even when the input and output sampling frequencies are the same. Obviously, this is cause for concern; we want perfect bit-for-bit accuracy between the original master recording and the signal being converted to analog. Asynchronous sample-rate conversion outputs a signal with perfect timing, but with samples of the wrong value. The technique converts errors in the time domain (jitter) to errors in the amplitude domain.

High-Resolution Digital Audio: Why 44.1kHz Sampling and 16-bit Quantization Aren't Enough

As explained in [Appendix C](#), the Nyquist theorem tells us that if the sampling frequency is more than twice as high as the highest audio frequency we want to encode, sampling is a theoretically lossless process. That is, it can perfectly encode frequencies at just under half the sampling rate. The Compact Disc's sampling rate of 44.1kHz was chosen because 44.1kHz is slightly higher than twice 20kHz, the highest frequency humans can reportedly hear. Why such an odd number? Early digital recorders were based on video formats such as 54" U-Matic tape, and 44.1kHz sampling happened to fit mathematically with the 30-frames-per-second rate of black-and-white video.

So if humans can hear only up to 20kHz at best, and CD's sampling rate of 44.1kHz can perfectly preserve a 20kHz sinewave, why does 96kHz sampling sound more open and transparent, and deliver greater transient detail?

When higher sampling rates became available in professional recorders in the mid-to-late 1990s, a debate raged over whether the sonic improvements were attributable to the preservation and reproduction of information higher in frequency than 20kHz, or to some other factor. That "other factor" turned out to be largely the relaxation in digital-filter requirements for higher sampling rates.

The digital filter in a CD player or outboard DAC removes all energy above half the sampling frequency. Unfortunately, digital filters introduce distortion in the time domain by smearing transients. Think of a snare drum that starts and stops very suddenly. Digital filters take some of that sudden energy and spread it out over time. Such distortion is also called ringing, because the filter "rings" like a bell when "struck" by the music signal. This ringing is figurative; a transient musical signal causes the filter to continue outputting a signal even after the musical transient has passed. The filter spreads some of that transient energy over time, smearing the signal. Oddly, this smearing occurs not only after the transient musical signal, but *before* it as well ([Fig.7-9](#)). The illustration shows an impulse (one sample at full amplitude) that is input to the filter (left), and the filter's output with ringing (right). The ringing that occurs before the musical transient is called pre-ringing. The steeper (i.e., more rapid) the filter's rolloff, the longer the period of time before and after the signal that its energy is smeared. In addition, steep filters introduce phase shift and amplitude-response ripples within

the passband. (The range of frequencies passed by the digital filter is called the passband. The frequency range stopped by the digital filter is called the stopband. The frequency range between the passband and the stopband is called the transition band.) This ringing, particularly pre-ringing, is largely responsible for the hardness and glare overlaying instrumental timbres, and the flattening of soundstage, that we associate with standard-resolution digital audio.




Fig.7-9 Steep digital filters smear music's transient detail.

The problems caused by having a steep digital filter right next to the audio passband are sometimes exacerbated by the low quality of the filter. It's possible to design filters that reduce these problems by employing greater mathematical precision, but they are too costly to implement except in the most expensive products. And because high-end companies often rely on makers of mass-market semiconductors for their digital filters, even expensive digital components don't always include high-performance digital filters. Moreover, the digital chain, from recording to playback, is a series of filters, each one degrading sound quality.

The steeper the filter, the greater the ringing and energy smearing, and the digital filters used to decode 44.1kHz audio are steep indeed; they must pass the entire audioband up to 20kHz, yet attenuate the signal by a whopping 120dB at 22.05kHz (half the sampling frequency). The passband is up to 20kHz, the transition band is between 20kHz and 22.05kHz, and the stopband is above 22.05kHz. For comparison, analog filters used in loudspeaker crossovers are much more gradual: in a typical two-way loudspeaker, the signal is attenuated by about 12dB over a full octave beginning at about 2.5kHz.

A filter designed for a 96kHz-sampled signal, however, can begin attenuating at a much higher frequency, and with a far gentler slope. The stopband for such a filter begins at 48kHz, the transition band is much wider, and the frequency at which the filter begins attenuating can be moved further away from the audioband. The result is less distortion—less amplitude-response ripple, less

phase shift, less energy smearing—within the audio passband, and, thus, better sound quality. For an academic treatise on this subject, see “A Suggested Explanation for (Some of) the Audible Differences Between High-Sample Rate and Conventional Sample Rate Audio Material,” by Mike Story of Data Conversion Systems (dCS), at www.cirlinca.com/include/aes97ny.pdf. Although the relaxation of digital filter requirements is certainly a factor in 96kHz’s higher sound quality, some respected researchers also suggest that even though we can’t hear sinewaves above 20kHz, we can nevertheless perceive the greater steepness—that is, speed—of transient musical information provided by ultra-wideband digital audio systems. The wider a system’s bandwidth, the faster the events (sudden changes in amplitude) it can preserve.

Consider a subwoofer, a very low-bandwidth device, reproducing the sound of a snare drum. The subwoofer removes only the very-fast-changing components of the signal (the steep transient of the attack), but reproduces the rest of the signal. Similarly, a 20kHz-bandwidth audio system filters out the steepest transients of musical signals, which perhaps introduces a kind of distortion. Some believe that the steeper transients reproduced by a wide-bandwidth digital audio system improve the sense of space, because the ear/brain uses such transient information to locate the sources of sounds.

Moreover, psychoacoustic research has revealed that subjects can identify a difference between a pure sinewave at 1kHz and the same sinewave with 25kHz added, though the same subjects are unable to hear the 25kHz sinewave on its own. This suggests that the ear/brain includes a mechanism for detecting “ultrasonic” sounds—sounds not directly audible to the human ear, but that are preserved in high-resolution audio.

The filters I’ve described are so-called linear-phase filters, which have been standard in CD players and DACs since the beginning of digital audio. Recently, however, designers have begun creating minimum-phase filters that sacrifice flat frequency response (they exhibit a slight upper-treble rolloff) for reduced ringing (better time-domain performance). These filters have become more popular as high-end designers increasingly create their own filters by writing custom software that controls a chip devoted to general-purpose digital signal processing (DSP). Often, the listener can select between filters by pushing a button on the front panel of the CD player or DAC.

In 2009, Meridian Audio introduced its so-called apodizing digital filter, which not only doesn’t introduce pre-ringing of its own, but can also remove pre-ringing added to the signal by processing earlier in the chain, primarily in the analog-to-digital converter. (The term *apodizing* comes from optics and radio

astronomy. Sharp edges at the boundaries of optical lenses or radio dishes create ripples in the diffraction pattern analogous to the ringing in digital filters. In radio telescopes, the contribution from the outer edge of the disc is attenuated to reduce this effect, a process called apodization.)

The pre-ringing that exists on nearly all the CDs in our collections is just as sonically detrimental as the pre-ringing added by the CD player's linear-phase filter. In fact, this pre-ringing is largely responsible for "CD sound"—that combination of flatness of soundstage and hardness of timbres. Meridian's apodizing filter, first used in their groundbreaking 808.2 Reference Signature CD player, greatly improved the sound of CDs, especially older discs that sounded particularly hard and flat.

Although apodizing filters are becoming increasingly common in a wide range of DACs, they don't actually remove the ringing energy. Rather, they shift the pre-ringing energy in time so that it occurs after rather than before the transient, making it less audible.

16 Bits, 20 Bits, 24 Bits

Although the causes of the superior sound quality delivered by increasing the sampling rate are debatable, the benefits of the increase in word length from 16 to 18, 20, and even 24 bits are not in dispute. As I mentioned earlier, word length is the number of bits used to encode the audio signal's amplitude at each sample. Assigning a number—called a word—to represent the audio signal's amplitude is called *quantization*. The word length determines the system's resolution, dynamic range, distortion, and signal/noise ratio. We also call the word length *resolution*, as in the phrase *20-bit resolution*. The more bits per word, the higher the resolution.

Think of attempting to communicate to another person the precise height of an object 1478.5mm high when you can use only one digit, four zeros, and a decimal point. The number that most precisely represents the object's height within this limitation is 1000.0, which tells the person that the height is greater than 500mm but less than 1501mm—not very high resolution. If you have two digits at your disposal, you can communicate the number 1500.00, which more precisely represents the object's height as being between 1451mm and 1549mm. With three digits, the precision, and thus the resolution, increases: 1480.0. With four digits, we get even closer to describing the object's height: 1479.0. Adding a fifth digit allows even finer resolution: 1478.5. The range of error is successively reduced with each additional digit. Similarly, the greater the number of bits in each digital word, the more precisely the analog signal's amplitude is encoded. Remember that digital audio attempts to describe an infinitely variable phenomenon—an analog audio waveform—with a finite number of discrete quantization words. Consequently, there will always be a difference between the analog waveform's precise amplitude and the quantization word assigned to represent that amplitude. The longer the quantization words, the more steps, and thus the finer the resolution.

The CD format uses 16-bit quantization, which provides 98dB of dynamic range. That is, it can encode a range of sounds almost as wide as that produced by a symphony orchestra. Although this number may sound high in theory, it's not enough in practice. Consider a sound engineer setting the recording levels for a recording of a symphony for full orchestra. He must adjust the signal driving the analog-to-digital converter so that, at the climax of the final movement, when every instrument is being played its loudest, the levels never exceed the full scale of digital recording; that is, the maximum loudness that can be encoded. At full

scale, all 16 bits are used. Now think about a quiet flute passage in the slow movement; that low signal level is encoded with perhaps only three or four bits, which offers very low resolution.

Consequently, the resolution of a digital audio signal isn't defined by the maximum number of bits available, but by the number of bits being used at any given moment. The lower the signal level, the lower the precision with which the signal is encoded. And if the engineer leaves 6dB of "headroom" (safety margin) to account for unexpectedly loud peaks, he's effectively throwing away one full bit of resolution, or 6dB of dynamic range.

The advent of 20- and 24-bit digital audio not only expands the dynamic range, but also increases the resolution of low-level detail. This low-level detail can be fine nuances of an instrument's timbre, which enhance the sense of realism. It can also be subtle spatial cues, such as discrete acoustic reflections and reverberation decay, which the ear interprets as a more convincing reproduction of the original recording venue.

Longer word lengths also contribute to better sound because the postproduction (mixing, equalization, signal processing) common in the recording or mastering studio can be performed with much greater mathematical precision. Moreover, any noise added by these processes is spread out over a wider bandwidth, which makes it less audible.

The combination of higher sampling rate and longer word length results in greatly improved sound quality. (A digital system's resolution can also be increased by adding to the signal a small amount of noise, called dither. See [Appendix C](#).)

The PCM format with the theoretically highest resolution is called *Digital Extreme Definition* (DXD). DXD is a PCM signal with a sampling frequency of 352.8kHz and a word length of 24 bits. DXD isn't a format for delivering music to consumers, but rather a professional format initially developed to edit Direct Stream Digital recordings when creating a master recording from which SACDs can be replicated. DSD's one-bit structure doesn't lend itself to signal processing, so it's converted to this very high-resolution PCM format for editing, then converted back to DSD for encoding on the SACD. Some upsamplers will convert lower-resolution PCM signals to DXD for decoding by a DXD-compatible DAC. Note that DXD's bit rate is a whopping 8.4672Mbps per channel.

A Caveat About High-Resolution Digital Audio

It's important to realize that a fast sampling rate such as 176.4kHz and a 24-bit word length don't guarantee good sound. Rather, those specifications simply provide the *opportunity* for higher-quality audio. Implementation is everything. A well-designed 44.1kHz/16-bit system will outperform a poorly conceived and executed "high-resolution" system.

Second, high-resolution analog-to-digital and digital-to-analog conversion is vastly more exacting than standard-resolution conversion. In a 16-bit system, the least significant bit of the 16-bit word represents one part in 65,536; in a 24-bit system, the least significant bit represents one part in more than 67 million. Many DACs don't even deliver actual 16-bit resolution, never mind full 24-bit resolution. High-res digital places demands not only on the DAC chip's low-level performance, but also on the analog output stage, clocking accuracy (jitter), power-supply purity, and every other subsystem. The temperature variations within components that might only marginally affect the performance of a 16-bit system can introduce gross errors in a 24-bit converter.

Third, the big numbers associated with high-resolution digital audio (96kHz, 176.4kHz, 24-bit) are often driven by marketing. The product might not actually deliver "96/24," but consumers often think that higher numbers automatically mean better sound. Although state-of-the-art high-res formats deliver 24 bits, and most digital-to-analog converters have 24 "rungs" in their resistor ladders, those last four bits generally lack precision sufficient to encode musical information in all but the best DACs. This is a limitation of the analog-to-digital and digital-to-analog converters. The additional 4 bits in a 24-bit DAC chip compared with a 20-bit chip are cynically called marketing bits; they serve no useful function in terms of sound quality. Although the DAC might have 24 "rungs" in its resistor-ladder network, those last four (or more) rungs are rarely capable of delivering real audio information.

Finally, a "high-res" datastream doesn't do you any good if the signal has previously been subjected to a standard-resolution A/D or sample-rate conversion. Once the music has been encoded at 44.1kHz (or 48kHz) and 16 bits anywhere in the chain, it has been irreparably compromised. This is an important consideration when buying "high-res" downloads. Have the engineers actually gone back and remastered the material from the original source elements in high-res, or have they simply upconverted the existing CD-quality datastream and called it high-res? And what about those albums tracked or mixed on the

48kHz/16-bit digital recorders prevalent in the late 1980s and early 1990s? This topic is discussed in more detail in the following chapter.

I raise these issues for two reasons: listeners should be aware that high-resolution specs are not a silver bullet for digital's flaws, and that high-res done right is rare.

Technical Addenda: Direct Stream Digital, HDCD, How a DAC Works

Direct Stream Digital (DSD) Encoding on SACD

Sony and Philips developed for SACD an entirely different method of digitally representing music, called Direct Stream Digital (DSD). Instead of sampling the musical waveform at a relatively slow 44,100 times per second and assigning a 16-bit binary number to represent the amplitude at each sample point (the parameters of CD), DSD samples the musical waveform at a lightning-fast 2,822,400 times per second. Each “sample” is, however, only one bit. If the waveform’s amplitude is increasing, a binary “one” is recorded. If the waveform’s amplitude is decreasing, a binary “zero” is recorded. An audio signal at full-scale positive is represented by all ones; a signal at full-scale negative is represented by all zeros. The zero point is represented by alternating ones and zeros, as shown in [Fig.7-10](#). The musical information is contained in the width of the pulses, which is why this type of encoding is sometimes called *pulse width modulation* (PWM). The data rate of DSD is exactly four times that of 44.1kHz/16-bit encoding ($44,100 \times 16 \times 4 = 2,822,400$ bits per second).

Interestingly, the pulse train generated by DSD encoding “looks” remarkably analog-like. That is, you can look at the pulse train and get an idea of what the analog waveform from the microphones looked like. Moreover, converting this pulse train to analog can be as simple as putting the pulses through an analog *low-pass filter*.




Fig.7-10 The DSD bitstream is a train of single-bit pulses. (Courtesy Sony Corporation)

In practice, the DSD-encoded signal is too noisy to convert directly to analog. Because this noise is distributed over a very wide bandwidth, a technique called *noise shaping* shifts the noise away from the audioband and up to a higher frequency, where it is inaudible. Although the amount of noise in the signal remains constant, the noise within the audioband is radically reduced. Fig.7-11 shows how noise shaping shifts noise away from the audioband.




Fig.7-11 Noise shaping shifts noise out of the audioband. (Courtesy Sony Corporation)

DSD can record frequencies of up to 100kHz—far higher than the 20kHz of conventional 44.1kHz-sampled digital audio. In addition, the signal/noise ratio (the difference between the loudest sound that can be recorded and the system’s noise level) is a whopping 120dB. This is contrasted with the 98dB signal/noise ratio of 16-bit quantization. Note, however, that DSD cannot achieve a signal/noise ratio of 120dB across its entire 100kHz bandwidth, but only through the audioband of 20kHz. That’s because the noise shaping described earlier

reduces the signal/noise ratio above the audioband. One of the benefits of DSD is the elimination of complex filters from the signal path. Conventional digital audio coding requires so-called decimation and oversampling filters, which introduce noise and errors. DSD is a much more direct method of digitally storing music.

Moreover, the DSD pulse train can be converted to 44.1kHz/16-bit form with fairly simple processing. The same DSD master tape can be the source for both the high-resolution and standard-resolution layers on an SACD. Interestingly, many of the sonic benefits of DSD are apparent in the 44.1kHz/16-bit signal. I've heard comparisons of music recorded on DSD, then downconverted to 44.1kHz/16-bit, and the same music originally recorded at 44.1kHz/16-bit. The downconverted DSD version sounded significantly better than the conventionally coded disc.

Before being recorded on an SACD master disc, the DSD bitstream is compressed so that it will take up less space on the disc. For example, a string of eight consecutive 0s may be coded as 8x0. (In practice, the compression system is vastly more complex.) Unlike "lossy" compression systems, which throw away portions of the signal considered to be "inaudible," the compression system used in SACD ensures perfect bit-for-bit accuracy to the source data by more efficiently coding the audio data. This compression system is what allows SACD to store both 2- and 6-channel representations of the music in its high-resolution layer.

A drawback of DSD is that all professional recording, mixing, mastering, and signal-processing equipment is based on pulse-code-modulation digital audio, not the 1-bit representation of DSD. As a result, many SACD titles have been recorded direct to two-track (or multitrack) DSD, with no mixing or processing after the recording. This approach maintains the greatest sonic purity, and is the preferred process for audiophile labels, but it isn't practical for the entire music industry. Sony has worked to provide recording and mastering engineers with DSD-based tools, but access to those tools is still limited.

DSD encoding provides the intriguing possibility of converting the pulse-width-modulated bitstream directly to analog in a *switching amplifier*. Rather than convert the DSD bitstream to analog and then amplify it conventionally, the DSD bitstream can be fed to a Class D switching amplifier, with the pulses turning on and off the amplifier's output transistors. An SACD player's DSD output is fed directly into the amplifier, with the amplifier's output stage functioning as the digital-to-analog converter. (Class D switching amplifiers are explained in [Chapter 5](#).)

High Definition Compatible Digital (HDCD)

The fundamental problem with CD sound is that the amount of digital data representing the music isn't nearly enough to correctly encode all the musical information we can hear. As explained earlier in this chapter, the CD's sampling frequency of 44.1kHz isn't fast enough, and 16 bits per sample are too few. The challenge, then, is to put more information about the musical signal through the compact disc's 44.1kHz/16-bit pipeline.

But how can you pack more information into a pipeline that's already filled to capacity? High Definition Compatible Digital (HDCD) uses some ingenious techniques to accomplish just that. The result is musical and technical performance that exceeds the limits of conventional 44.1kHz/16-bit digital audio. The process was invented by Keith Johnson and Michael Pflaumer, who, with Michael Ritter, formed Pacific Microsonics Inc. to develop and market HDCD. In 2000, Pacific Microsonics was acquired by Microsoft, which has a keen interest in the art of transmitting more information in fewer bits. Many discs are encoded with HDCD, players with HDCD decoding are still available, and HDCD offers, in my view, a significant improvement in CD performance.

HDCD is used primarily at the mastering stage of the recording chain. After an album has been recorded and mixed, the producer and artist take the final tape to a mastering studio to generate the digital master tape from which CDs are made. The mastering engineer, artist, and producer decide if they want to use the HDCD process.

Recordings encoded with HDCD can be decoded and played back on any CD player, with some improvement in fidelity. But if your CD player is equipped with an HDCD decoder, HDCD recordings can be reproduced with much better sound quality than is possible with conventional CDs. Here's how it works.

The HDCD encoder takes in an analog audio signal (either from analog tape or a live microphone feed) and converts it to digital with a very fast sampling rate (176.4kHz) and 24-bit resolution. The amount of data generated is far too great to store on CD, so the HDCD encoder analyzes the high-resolution signal to determine what musically significant information will be lost when that signal is converted into a 44.1kHz/16-bit signal for storage on CD. The encoder then outputs a 44.1kHz/16-bit digital audio signal, but with a control code hidden in the 16th bit. This hidden and inaudible control code tells the HDCD decoder in the CD player to perform certain processes on the audio signal to restore some of

the information present in the original high-res digital signal. (Note that you don't give up one bit of resolution with an HDCD recording; the control code is present only 3–5% of the time.)

When decoded, the small amount of information in the control code can produce large changes in the audio signal. That's because the control code is a kind of shorthand that the decoder recognizes. Here's an analogy: Say we want to transmit Beethoven's Ninth Symphony from one place to another. Rather than transmit the entire symphony, we instead send just a few bits that tell another device to output Beethoven's Ninth Symphony. That's how HDCD works: instead of transmitting the musical information that would normally be lost in conventional coding, the control code transmits instructions to the decoder telling how to restore that information. Because so much of HDCD's "intelligence" is in the decoder, a few bits in the control code accomplish the work of many.

Although HDCD uses many techniques to improve sound quality, let's look at how HDCD overcomes the limitations of the CD's 20kHz bandwidth.

While it's true that humans can hear pure tones up to only 20kHz (and only if you're 16 years old; at age 40, you're lucky to hear 15kHz), that doesn't mean there's no perceptible musical information above 20kHz. Very high frequencies imply a steep attack; the ear/brain can reportedly detect the difference in steepness of a signal with a bandwidth of 20kHz and one of 40kHz, according to HDCD's inventors (and many others).

The bandwidth of an audio system implies how quickly it can respond to musical signals. Earlier in this chapter I described a subwoofer, a low-bandwidth device, trying to reproduce the attack of a snare drum. The woofer simply can't move fast enough to accurately reproduce the steep transient of the drum's waveform.

That's exactly what happens when we record music with a bandwidth of only 20kHz. The 20kHz-bandwidth digital audio system simply can't encode the parts of the music that quickly change in intensity. The musical result is a closing-in of the sound and a reduced ability to pinpoint the positions of instruments. Experiments have shown that our ears and brain use that steep transient information as directional cues. Moreover, limiting the speed of musical signals changes their timbre. For example, the component of an oboe's sound that includes the tiny "clicks" of the vibrating reeds disappears. We don't hear this as distortion, but as a dilution of realism. This is one reason that a live oboe has a vitality and presence we don't hear from a standard-resolution digital

reproduction of an oboe. Because the CD's bandwidth isn't wide enough, instrumental timbres are changed, along with their directional cues.

Remember that control code hidden in the 16th bit of an HDCD-encoded recording? Part of that control code tells the decoder in your CD player that the attack transient of a particular signal was slowed; the decoder then restores the original steepness of the transient. This is how HDCD can reproduce signals whose speed implies a bandwidth greater than 20kHz within the limitations of the CD's 44.1kHz sampling frequency.

This is just one of the many techniques HDCD uses to squeeze more music through the 44.1kHz/16-bit pipeline. Similar techniques address the dynamic-range limitations of 16-bit coding. The system is extremely sophisticated in practice, and reflects a decade of research into human hearing, recording, electronic design, and digital signal processing. Note that HDCD-encoded CDs are encoded at a level 6dB lower than conventional CDs to allow for the greater dynamic headroom restored by the HDCD decoder.

The first HDCD-compatible products used the Pacific Microsonics PMD100 HDCD decoder, which included an 8x-oversampling digital filter that is integral to HDCD decoding. This filter also operates on conventionally encoded discs, which makes all recordings sound better, not only HDCD-encoded ones. That's because the HDCD chip's integral digital filter operates with a higher mathematical precision than previous filters, and employs some tricks to improve sound quality. Today, HDCD decoding is more likely to be incorporated into a DSP chip that performs many other functions. Microsoft sells professional HDCD encoders and licenses HDCD decoding to other companies. For example, Motorola, Analog Devices, and Sanyo manufacture audio-processing chips that include HDCD decoding, and pay Microsoft a royalty on every chip sold. This licensing agreement is similar to that used by Dolby Laboratories in licensing its ubiquitous Dolby B noise-reduction circuit and Dolby Digital decoding technology.

In my view, HDCD is a significant advancement in digital audio sound quality. I've had the opportunity to compare an original analog master tape with an HDCD-encoded/decoded version, along with conventional 44.1kHz/16-bit coding. I've also heard comparisons of a live microphone feed (and the instrument in the hall) against state-of-the-art analog tape and HDCD. My experience suggests that HDCD improves sound quality in these ways: with a finer resolution of detail, more accurate rendering of timbre, smoother and less synthetic treble, wider dynamic range, clearer articulation of loud and complex passages, bigger soundstage, tighter image focus, more front-to-back layering,

more bloom around instrumental outlines, and a greater projection of quiet instruments in the presence of louder instruments.

Interestingly, many HDCD-encoded CDs lack any indication on the disc or packaging that the disc is so encoded. The Pacific Microsonics analog-to-digital converter and HDCD encoder are used in a large number of mastering rooms primarily because they're generally considered the best-sounding A/D converters extant. Albums mastered through a Pacific Microsonics A/D are often HDCD encoded by default, and the record company simply doesn't know enough to indicate on the packaging that the disc has been HDCD encoded. It's a pleasant surprise to buy a CD and discover, when you play the disc for the first time and see the HDCD indicator on the CD player or DAC illuminate, that the disc has been HDCD encoded.

How a DAC Works

The rest of this chapter delves into the inner workings of DACs. You don't need to know this information to enjoy music in your home; I've included it for the more technically minded—those who want to understand what goes on inside those sculpted metal boxes in their equipment racks. If you like, simply skip ahead to the beginning of the next chapter.

Fig.7-12 shows a block diagram of an outboard DAC. The main components are a power supply, input receiver, digital filter, digital-to-analog conversion stage, current-to-voltage converter, and analog output stage.

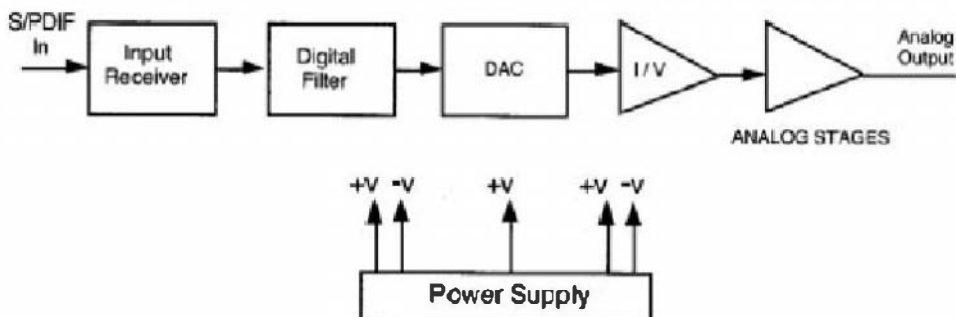


Fig.7-12 A DAC's main subsections are the input receiver, digital filter, digital-to-analog converter chip, current-to-voltage converter, and analog output stage.

The input receiver takes the S/PDIF signal from a digital source and converts its serial datastream into raw digital audio data. It also generates a clock signal based on the clock embedded in the datastream. A phase locked loop (PLL) compares the input signal frequency (the clock) with a reference frequency (usually generated by a crystal oscillator), and creates a new clock locked to the incoming datastream's clock. This so-called "recovered" clock becomes the master clock for the processor. The input receiver is a significant source of clock jitter, and can have a large effect on how the processor sounds. Recent attempts to minimize jitter created by the input receiver include dual PLLs, custom low-jitter input receivers, and external clocks.

Before we move on to the next step in digital-to-analog conversion, let's look at what happens when we convert a digital signal to analog. The discrete samples that represent the analog signal don't form a smooth, continuous waveform; they instead create a stairstep waveform. The energy in the stairstep waveform's sharp

edges appears as images of the signal centered at multiples of the sampling frequency (Fig.7-13). These images must be removed by an analog low-pass filter.

In the first Sony players, these images were removed and the staircase waveform smoothed by a brickwall analog filter (top drawing, Fig.7-13). This filter was extremely steep and close to the audioband, and resulted in phase shift, ripple, and other audible artifacts.




Fig.7-13 Spurious images are created at multiples of the sampling frequency (top). An over-sampling digital filter shifts the first image to a higher frequency that is more easily filtered. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

Today's CD players and DACs use an oversampling digital filter that increases the sampling frequency by (typically) eight times by interpolating seven new samples for each audio sample. This raises the sampling frequency to 352.8kHz ($44.1\text{kHz} \times 8$), which pushes the first image to 352.8kHz, where it can be filtered by a much gentler and better-sounding analog filter well away from the audioband (lower drawing, Fig.7-13). Again, the spurious images at multiples of the sampling frequency still appear, but are now far outside the audioband. Contrary to some beliefs, oversampling's sonic benefits are not the result of adding new information between the existing samples, or from "smoothing out" the waveform with more samples. (The interpolated samples are actually all zeros, which contain no information.) Moreover, oversampling doesn't increase the audio bandwidth; an 8x-over-sampling CD player still has a bandwidth of 20kHz. The sonic advantage is achieved by replacing the brickwall filter with a gentler analog output filter.

In addition to increasing the effective sampling rate, oversampling digital filters increase the output word length from 16 bits to, typically, 24 bits. Note that no new information is generated by the filter when outputting longer words. Instead, the longer words reduce errors resulting from the filter's internal calculations. The audio data being filtered may temporarily have very long word lengths—perhaps 30 bits—that need to be reduced before being output by the filter. Because simply truncating (cutting off) all but the 16 needed bits results in audible distortion, some filters use redithering techniques (see [Appendix C](#)) to reduce round-off errors. These techniques attempt only to compensate for errors introduced in the filter; they cannot improve incoming data.

Custom Digital Filters

Manufacturers of CD players and DACs have two choices: buy an “off-the-shelf” filter chip, or create a custom filter from general-purpose digital signal processing (DSP) chips. The designer of the custom filter must write the software that controls the DSP chips—an expensive and time-consuming task. Consequently, custom filters are quite a bit more expensive, but give the DAC designer greater creative control over the product’s sound. In addition, custom digital filters can run faster than the typical 8x oversampling of chip filters.

Proponents of this approach believe that their custom filtering software is better than that found in conventional integrated circuit chips. Specifically, most custom digital filters are optimized for time-domain rather than frequency-domain performance. As described earlier in this chapter, some digital filters accept a rolloff at the top of the audioband (typically, -3dB at 20kHz) to reduce filter ringing.

Digital-to-Analog Conversion

After digital filtering, the oversampled data are input to the digital-to-analog converter (DAC). The DAC converts digital data at its input to an analog output signal. The most common type of DAC is the multibit “ladder” or “R/2R” DAC, both names describing the same operating principle. Multibit ladder DACs have “rungs” of resistors, each rung corresponding to one bit in the digital code. A 16-bit DAC will thus have 16 resistor rungs, which allow 65,536 (2^{16}) possible input codes. The resistor value of each rung is twice the value of the adjacent resistor, thus the name “R/2R” ladder. The input data act as switches on the rungs; a binary “1” closes the switch at that particular bit and allows current to flow through the resistor, while a binary “0” opens the switch and no current flows. The DAC’s output current is thus determined by the digital input code. (An R/2R resistor-ladder DAC is shown in Fig.7-14.)




Fig.7-14 An eight-bit R/2R ladder DAC (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

Each rung has an effective resistor value half that of the adjacent rung; the resistor determines how much “weight” the bit has by determining how much current flows through the switch. The least significant bit (LSB) has the lowest weight, the most significant bit (MSB) the highest. Each bit upward from the LSB should produce a precise doubling of current, in the progression 1, 2, 4, 8, 16, 32, 64, 128, 256, and so on, up to 524,288 in a 20-bit DAC. The LSB has a value of 1, the MSB a value of 524,288 (in a 20-bit DAC). When the MSB value of 524,288 is combined with the 19 lower bits, a 20-bit DAC can accommodate up to 1,048,576 input codes.

The precision of these resistor values is crucial: Any deviation from a doubling of resistance causes that bit to have more or less value than it should. These bit-weighting errors result in poor linearity. That is, the analog output level doesn’t track the digital input code. A perfectly linear DAC will produce an analog output at exactly the same level as the digital input. For example, if the DAC is driven by the code representing a -90dB dithered sinewave, the analog output should be exactly 90dB below full scale. A nonlinear DAC would produce an output level of perhaps -93dB or -86dB. Linearity errors generally occur at very low signal levels.

Delta-Sigma DACs

An alternative to the multibit resistor-ladder DAC is the *delta-sigma DAC*, also called a 1-bit DAC. These converters are sometimes called by their trade names: Bitstream (Philips), MASH (Matsushita's process, developed by Nippon Telephone and Telegraph), and PEM (Pulse Edge Modulation, developed by JVC). All of these converters work on the same principle: Instead of using a resistor ladder with different bit weightings to convert a binary code into analog, 1-bit DACs operate in only two states: zero and one. The 1-bit code is a series of varying-length pulses with a single amplitude. The DSD encoding system used in SACD is delta-sigma modulation.

The single-bit code is of a high enough frequency (typically between 64x and 256x oversampling) that the audio signal can be reconstructed from these two logic states with a switched-capacitor network. Consequently, delta-sigma DACs don't require the amplitude precision of ladder DACs. Delta-sigma DACs trade resolution in amplitude for resolution in time. They have inherently good linearity, and don't require a current-to-voltage converter. Because delta-sigma DACs operate at such high frequencies, noise shaping is used to shift quantization noise from the audioband to just above the audioband.

Multibit resistor-ladder DACs and 1-bit DACs are being replaced by chips that combines the best aspects of both approaches. These DACs, called multibit delta-sigma DACs, are increasingly popular. The multibit word (typically 16 to 24-bits long) is converted into a five-bit code, but at a higher sampling frequency. No information is lost in this conversion. Because the DAC has only five rungs on its R/2R ladder, the resistor values needn't be as precise as in a 24-bit resistor ladder.

The British company Data Conversion Systems (dCS) has developed their own novel topology that is a hybrid of multibit and delta-sigma techniques. Their so-called Ring DAC is a five-rung resistor-ladder DAC in which each resistor value is identical. A modulator before the DAC converts the PCM or DSD code into a five-bit form running at 2.8224MHz, four times the data rate of CD. Using identical resistor values solves the problem of multibit DACs in which any variation from a doubling of resistor value between adjacent rungs introduces linearity error. Because even nominally identical resistors have some differences due to manufacturing tolerances, the signal is shifted (passed around the "ring") between resistors so that any errors are randomized and converted into a small amount of white noise. [Fig.7-15](#) is a graphic comparison of sampling rates and

quantization levels between 44.1kHz/16-bit PCM, high-resolution PCM, the five-bit Ring DAC format, and delta-sigma (DSD) conversion.

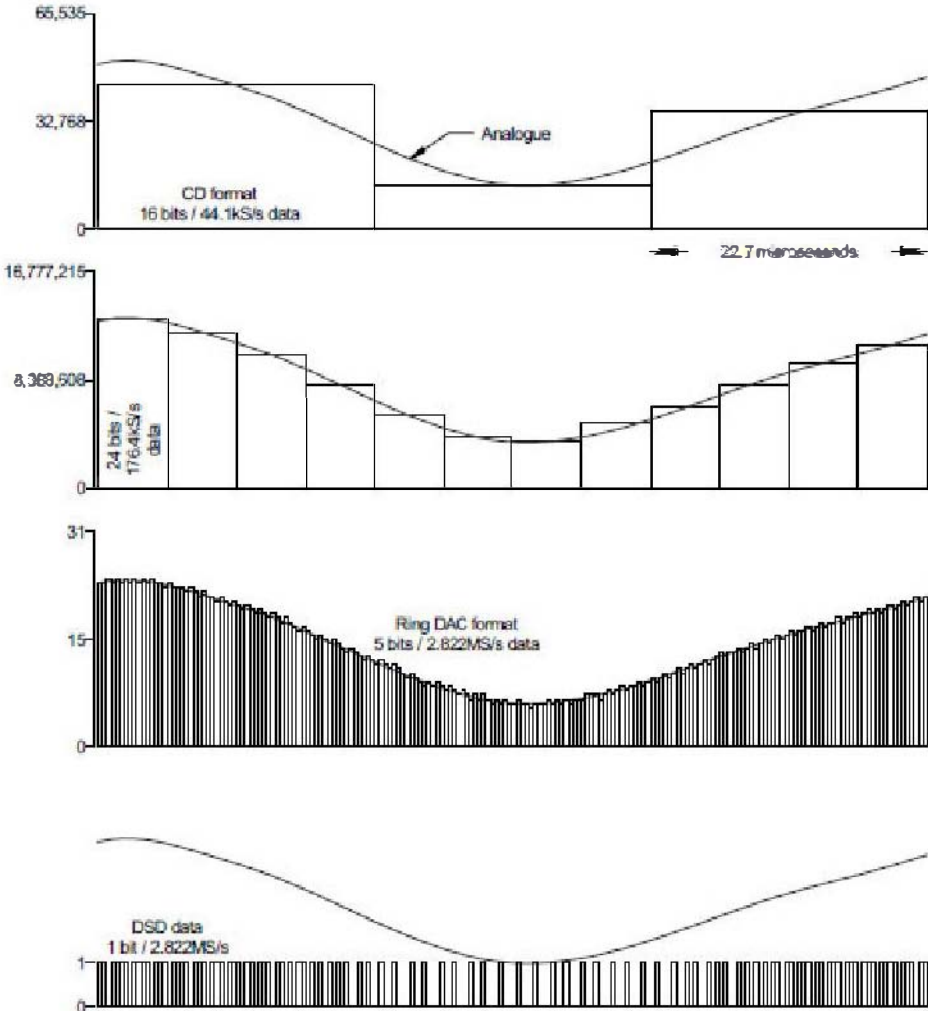


Fig.7-15 A graphic comparison of sampling rates and quantization levels. This conveys at a glance the respective data densities of each coding scheme. (Courtesy dCS)

Analog Stages

The output of a multibit resistor-ladder DAC is a series of currents at varying levels. This current output must be converted into a voltage that can be amplified by the DAC's analog output stage. This is the next step after the DAC, called current-to-voltage conversion, or I/V conversion (I is the electronic symbol for current, V the symbol for voltage). The I/V converter converts the DAC's current output to a voltage. For I/V conversion, virtually all DACs use an operational amplifier (op-amp), an integrated circuit that takes the place of a circuit using separate transistors. In addition to being easier and cheaper to implement, an op-amp is generally faster than a discrete (separate transistor) I/V stage. This is because the circuitry is on a single silicon chip rather than spread out in multiple, discrete components on a circuit board.

After I/V conversion, the analog voltage is put through an analog output stage. This circuit has several functions. First, it can provide gain to bring the final output level to the standard of 2V RMS (with a full-scale digital input signal). Second, the output stage incorporates an output filter to remove high-frequency noise above the audioband. This filter is often a gentle first-order type. The output filter can be an elaborate active circuit, or as simple as a capacitor put across the analog output jacks. Finally, the output stage acts as a buffer between the I/V converter and the final analog output. The output stage generally has a low output impedance, and is designed to drive cables and a preamplifier. The analog output stage can be based on an op-amp or, in more expensive products, a discrete analog stage composed of separate transistors. A direct-coupled output stage has no capacitors in the signal path to block direct current (DC) from getting into the audio signal. Most direct-coupled output stages use a DC servo to prevent DC from appearing on the analog output jacks. A DC servo is a circuit that monitors the DC level at the amplifier's output and makes adjustments to reduce the amount of DC.

Balanced Outputs

Designers can use one of two very different techniques to create a balanced analog output signal. The easiest way is to simply put a phase splitter in the DAC's analog output stage. A phase splitter is an electronic circuit that generates a second signal of opposite polarity to the input signal. These two signals are then connected to an XLR jack on the DAC's rear panel. This technique, shown in block form in [Fig.7-16](#), is the inexpensive and easy way to create a balanced signal.

Some disc players and DACs, however, use a much more elaborate scheme. Rather than create the balanced signal in the analog domain with an additional active device (the phase splitter), these products create a balanced signal in the digital domain before digital-to-analog conversion. The digital datastreams for the left and right audio channels are split into two balanced signals (left +, left -, right +, right -), which are separately processed and converted to analog. This technique requires nearly double the circuitry, with four digital-to-analog converter chips and four sets of analog output electronics. A digitally balanced DAC is shown in block form in [Fig.7-17](#).

This method, though much more expensive than a simple phase splitter in the analog stage, has many advantages. When the balanced signal is eventually summed (in the preamplifier or power amplifier), any noise or distortion common to both channels will be rejected. This phenomenon, called common-mode rejection, is described in [Chapter 10](#), "Cables and Interconnects." Any artifacts or spurious signals introduced by the DACs will thus cancel each other out. Moreover, balanced DACs yield a 3dB increase in signal/noise ratio. Finally, a fully balanced DAC doesn't add another active device (the phase splitter) to the signal path.

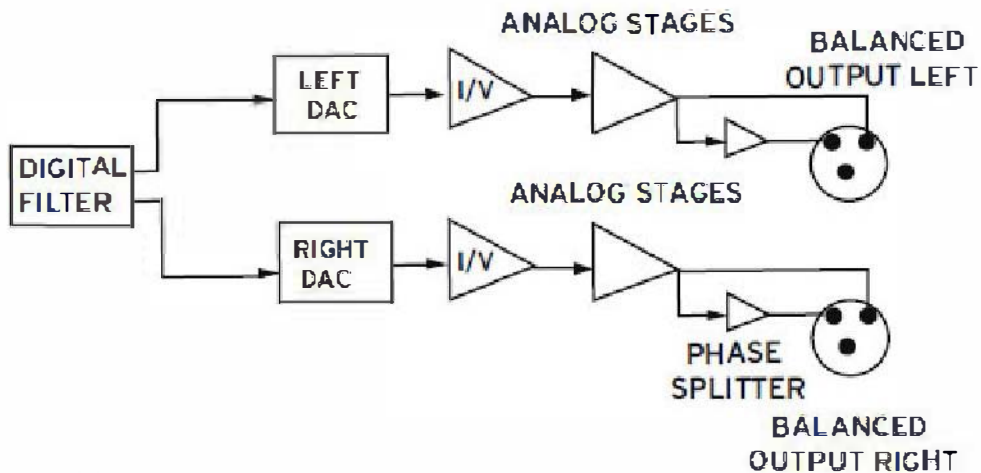


Fig.7-16 A “balanced” DAC can be created simply by adding a phase splitter to the analog output stage.

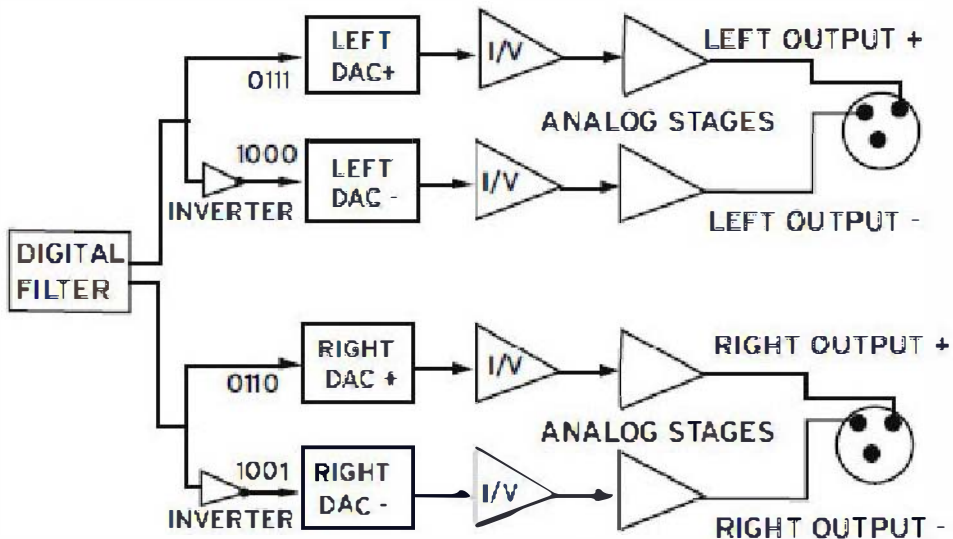


Fig.7-17 A truly balanced DAC splits the signal into “+” and “-” phases in the digital domain, and then converts the signal to analog with two converters per channel.

A final consideration with fully balanced DACs is how the unbalanced signal is derived. Some products merely connect half the balanced circuit (the left- and right-channel + signals) to the left and right rear-panel RCA jacks. Others sum the + and - signals with a differential amplifier (a circuit that converts a balanced to an unbalanced signal) to drive the unbalanced RCA output jacks. The latter approach provides the benefits of balanced DACs even for those users who don't use the balanced XLR outputs. Conversely, taking half the balanced signal at the

unbalanced output obviates the advantages of fully balanced DACs when using the unbalanced outputs. A comparison of these two techniques is shown in Figs.7-18a and 7-18b on the following page.

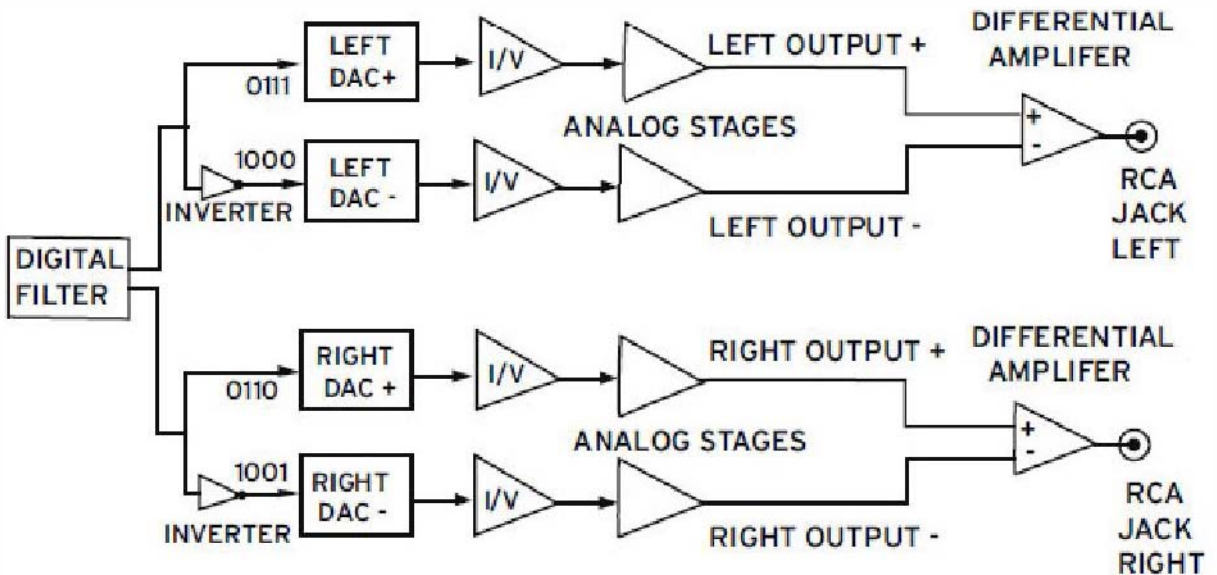


Fig.7-18a A balanced DAC's unbalanced output can be generated by combining the "+" and "-" phases of the balanced signal.

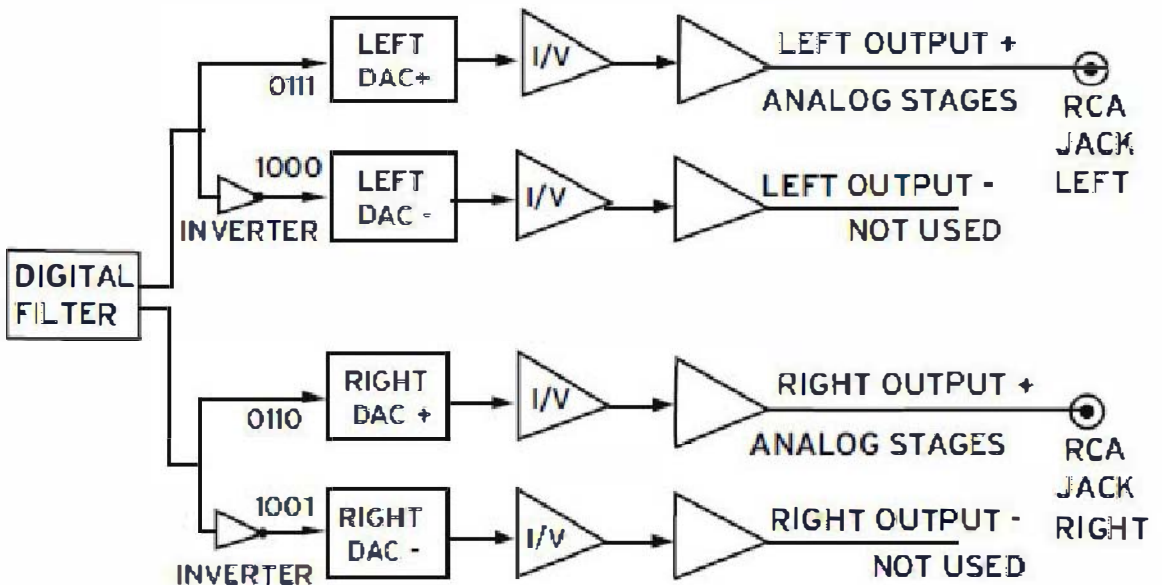


Fig.7-18b A balanced DAC's unbalanced output can be generated simply by using one phase of the balanced signal.

8 Digital Source Components, Part Two: Music Servers

The term *music server* once meant a personal computer configured to store and manage a music library, or a ready-made turnkey system that performed the same function. But the past few years have seen an explosion of entirely new product categories that offer a rich, varied way of accessing and playing *file-based* music. This term refers to acquiring, storing, organizing, managing, and playing music recordings that have been encoded as computer files. File-based music encompasses the products we think of as music servers, but also a new generation of devices that open up novel possibilities for accessing music in the home or on the go. Moreover, many of these innovative products are small, sleek, “smart” devices with fresh appearances that bear little similarity to the big, heavy, black boxes of traditional audio components. In this chapter we’ll look at the entire spectrum of options available to the music lover for accessing music stored as digital files.

If you’re hesitant about making the leap to a music server, the following description may make your decision easier.

Imagine sitting in your listening chair and deciding that you want to hear a particular piece of music. You get up, walk over to racks full of CDs, turn your head sideways to read the labels, and scan the jewel-box spines for the desired disc. You then take a jewel box from the rack, remove its CD, put it in a disc-transport mechanism, start the disc playing, and return to your chair.

Now, imagine opening an app on your tablet or mobile device and seeing your entire music collection displayed as a series of album covers. The titles can be presented alphabetically—by genre or composer or conductor or orchestra or, in many cases, by which musician(s) plays on the album. You touch the image of the album cover and the music instantly begins to play.

Welcome to the world of music servers. But what I've described is only the beginning of how a music server can transform your relationship to your music library. Say that, in the recording just selected, you found yourself greatly appreciating the drumming of Peter Erskine. On what other records in your collection does Erskine play? A couple of finger taps on the touchscreen display show you all the titles featuring this virtuoso drummer, ready for instant listening. Or consider a scenario we've all found ourselves in: We know what kind of music we're in the mood to hear, but not the specific title. Choose a title on the touchscreen that's in the genre you're interested in hearing, and the music server rearranges the album art on the screen to show you other titles you might want to hear. Or simply choose a genre and allow the server to surprise you by playing music from that genre for as long a time as you specify (a feature that's great for parties).

The server is not only a storehouse for your music collection, but a portal for accessing new music. Music-download websites are like huge music stores that never close. In addition to buying downloads to store and play on your server, you can join a subscription service and have access to just about any music you want, and as much as you want, for a flat monthly fee. In fact, distinctions of where the music comes from are increasingly blurred—you have instant access to a vast music library from the listening seat, all organized in a way that lets you find exactly what you want. These technologies offer unprecedented opportunities to make musical discoveries. Anyone who has lived with a music server quickly finds it indispensable. In fact, the concept of music being stored on individual pieces of plastic that must be organized, shelved, found, and inserted into a machine, all by hand, will one day seem a quaint anachronism—and, to many listeners, already does.

Music Server Sound Quality: Better than Disc

A music server not only allows you to explore your music collection in new and rewarding ways, it also has the potential of delivering better sound than is available from CDs. Music read from a hard-disk drive can sound better than when read from an optical disc. That's correct—music ripped from a CD to a music server can offer greater fidelity on playback from the server than from the original disc itself, when decoded by the same digital-to-analog converter.

How can a copy sound better than the original? Consider what goes on during playback of a Compact Disc: a laser reads the data off a disc spinning at 200–500rpm, performing error correction on that data on the fly. The optical system has one pass to get it right. Moreover, the crucial clock that controls the process of digital-to-analog conversion is tied to the clock that controls the spinning disc.

Now consider what happens when you *rip* that data to a hard-disk drive on a server. If errors are detected, the disc-ripping software can read areas of the disc repeatedly, moving on only when the data are perfectly accurate. Those data are stored on a hard drive and treated like any other file by the computer. Reading those data off a hard drive is a cakewalk compared with the challenge of giving a laser only one pass to correctly track and retrieve all of the data from a spinning, wobbling, eccentric optical disc. The computer can read the data into a short memory buffer, then clock it out of that buffer with high precision. Moreover, there's evidence that the physical quality of the disc itself—specifically in the pit and land structures in which the audio data are encoded—affects sound quality, a factor obviated by hard-disk storage.

I've compared the sound of state-of-the-art CD transports to music servers playing the same music through the same digital-to-analog converter and digital cable. I ripped the CD to the server, put the CD in the transport, then compared them by playing both and switching between them. The sound from the server was smoother in the treble, with a more delicate rendering of treble detail. The soundstages of music sourced from the server were also improved—more spacious, greater image focus, increased depth, and an overall more convincing impression of real instruments being played in a real space in an actual hall. There was also a greater sense of ease and musical involvement. In short, transferring your CDs to a server will improve the sound of your CD library—provided that the server is configured and set up correctly. There's one other

caveat: The DAC's USB interface must be designed to the highest standard for the music server to sound better than disc, and the USB cable must be of high quality.

Downloads

A music server can deliver better sound than CD also because of its ability to play high-resolution music files. An increasing number of music titles in high-resolution sound are available via Internet download. The combination of a music server, high-quality digital-to-analog converter, and true high-res source material is revelatory. In fact, high-res digital done right bears little resemblance to most CD sound.

More and more high-resolution music is available from a variety of audiophile-oriented websites. Although the iTunes store has increased the maximum resolution of the files they sell, that music is still sonically compromised, and a long way from high resolution, or even CD resolution. That's why you'll need to buy downloads from one of the audiophile-oriented websites.

Download Caveats

Many high-res titles sound spectacular, and far better than the CD version of the title. Yet others offer only mediocre quality that's no better than CD, even though the file is delivered at 96kHz/24-bit. In some cases, the "high-res" file is simply a standard-resolution (44.1kHz/16-bit) signal upconverted to 96/24. Moreover, even if the file began life as a true high-res signal, subsequent signal processing can destroy the qualities we seek in high-res digital audio. If the file was ever converted to 44.1kHz, then the damage has already been done—nothing can return it to any former glory it may have had.

In addition, some recordings from the 1980s through the late 1990s were recorded on professional digital recorders that recorded at a sampling rate of 48kHz. If that's the case, downloading a version at 96kHz offers no sonic advantage. The consumer has the right to know the provenance of the claimed high-resolution file offered for sale. The best-sounding high-res files were recorded in high resolution at the session, or transferred at high res from original analog master tapes, then delivered to the consumer in a form that maintains perfect bit-for-bit accuracy to the master digital file.

A somewhat related issue is that high-res versions of recordings are sometimes subjected to the same detrimental signal processing used for the standard-resolution release. It's become common practice for mastering engineers of pop and rock recordings to severely compress a recording's dynamic range so that it sounds louder when played on the radio. Record producers feel the need to make their product more competitive, in terms of sheer volume, when played just before or after other recordings. The result of these so-called loudness wars is, in my view, an unnecessary and tragic destruction of much of music's range of expression. Highly compressed recordings sound uninvolved, flat, and hard, and lack rhythmic drive.

Unfortunately, virtually all CD titles that have been mercilessly compressed are also compressed in their high-res versions. For a variety of reasons, the high-res version will still sound better than the CD, but it won't deliver the full potential of which high-res digital audio is capable.

Some download sites offer titles in a wide range of file formats and resolutions, such as 96kHz/24-bit FLAC or 192kHz/24-bit WAV. These file formats are explained later in this chapter. It's also possible to buy downloads in the Direct Stream Digital (DSD) format, also described in this chapter.

Turnkey Servers vs. Computer-Based Servers

Some music servers can be called *turnkey*; i.e., you take it out of the box and start loading your music into it (or have the music loaded into it for you—more about this later). Other than that, a turnkey server requires no effort or expertise on the user's part.

The alternative is to create your own server from a personal computer. It's been said, only half jokingly, that the turnkey music server is for the ultrarich, and the *do-it-yourself* (DIY) computer-based server is for the ultrageeky. That's an exaggeration, but it highlights the fundamental differences between the two approaches: turnkey servers are expensive, and computer-based servers can be technically challenging. Let's first look at the spectrum of turnkey music servers, and later in this chapter consider the DIY server.

The Turnkey Music Server

The turnkey music server is, at its heart, a personal computer, but that fact is transparent to the user. In a well-designed server, the familiar computer user interface is replaced by a control system skillfully tailored to perfectly suit the system to the organization of a music library. In fact, much of what you pay for in a turnkey system is the development of the software that makes possible the system's ease of use. You don't need to worry about installing and configuring music-management software, going into the computer's operating system and making sure that all the settings are correct, and spending time on the phone with technical support trying to get the system to work correctly.

Most turnkey systems, such as the one shown in [Fig.8-1](#), consist of a touchscreen display along with another box containing the computer hardware: the hard-disk drives on which the music is stored. The turnkey system can forgo a dedicated touchscreen in favor of a phone or tablet app that provides the same functions.

When you turn on the system, your entire CD library is displayed, alphabetically by artist, as a series of album-cover images. You can manually scroll through the pages of covers, or jump to a specific artist's name by touching the letters on an alphanumeric display. Touching an album cover brings up a screen showing information about the record, including the list of tracks. You can select specific tracks to add to the playlist, or add the entire album. You can add as many tracks or albums to the playlist as you like.

This method of selecting music to play is far more convenient than finding a CD on a shelf, taking it out of the case, and putting it in a CD player. Most of us have music libraries that are far too big to scroll through, and we often approach a listening session not knowing specifically what we want to hear. If that's the case, you can select a musical genre and see just the titles within that genre. Other features include the ability to search for an artist, album, or song by name. Your library can also be displayed by file format and resolution—all FLAC or 192kHz/24-bit files, for example. Music servers require an always-on Internet connection. When you insert a CD into the server's CD-ripping drive, the system downloads the disc's album art and metadata from the Internet (specifically, from AllMedia). In addition to the usual metadata (album credits, playing times, track titles, release date, etc.), some turnkey servers offer additional information about the album or artist, described later under "Metadata."



Fig.8-1 A turnkey touchscreen music server. (Courtesy Meridian Audio)

There are several ways to get music into a turnkey system. The first, obviously, is to manually load each CD in your collection into the server's transport. But at about eight minutes per disc, importing a large library soon becomes tedious. Some server companies offer the option of shipping the server already loaded with music you specify such as a package of essential rock records, or the entire output from a specific record label over a specific time frame (all Prestige Records releases from the 1960s, for example). The cost is a little less than what it would cost to buy the music on CD and rip it yourself. You can also send your CDs and a hard-disk drive to a CD-ripping company that will rip your CDs to the drive. Whatever your solution, getting lots of music into the server is paramount—the more music loaded, the better the experience, and the greater the benefits of the server's sophisticated features.

The server can not only feed your main audio system, but also stream digital audio to other rooms of the house wirelessly (WiFi) or through wired connections (Ethernet).

Another Type of Turnkey Music Server

The turnkey system I've just described is as an example of a product at one extreme of the spectrum of components that deliver file-based music—a product that requires zero user expertise and offers the ultimate user interface. Another class of products offers the virtues of file-based music in a different sort of package. Let's look at a few of these.

Fig.8-2, the N500 from Astell&Kern, is a complete music system that provides, in a single component, a CD ripper, hard-disk storage, DAC, wired or wireless network streamer, power amplification, and a user interface with a 7" touchscreen. The touchscreen folds down into the top of the chassis when not in use, or when the system is controlled remotely via an iOS or Android app. When connected to a home network, the N500 will stream music from any PC, Mac, or *network-attached storage* (NAS) drive on the network. The system is compatible with all high-resolution formats, including double-DSD (DSD 128). When the N500 is used with Astell&Kern's high-end portable player (described in [Chapter 13](#), "Personal Audio"), you can wirelessly stream content between the home server and the portable player in either direction.



Fig.8-2 Music servers are increasingly taking on new capabilities and shapes. (Courtesy Astell&Kern)

A more specialized component is the product shown in Fig.8-3. This device, the Aries from the company Auralic, turns a DAC that supports USB Class 2 Audio (not to be confused with USB 2—both are explained in the previous chapter) into a *wireless streamer* that can decode virtually any audio format. The product accepts digital inputs from a wide range of sources (including streaming services, outboard devices, WiFi, Ethernet, and NAS and USB drives) at any sampling rate including DSD, double DSD, and DXD, then drives your DAC via the Aries's output. The product is supplied with an app (iOS or Android) for browsing and managing your music library. The output, which has been clocked with high precision, appears on USB, AES/EBU, coaxial S/PDIF, and TosLink. This category of product also offers seamless integration with streaming services such as Tidal (formerly called WiMP in Europe) and Qobuz: you can access music from these services via the component's tablet or phone user interface.



Fig.8-3 Components for playing file-based music look very different from the large black boxes of traditional audio products. (Courtesy Auralic)

Another example is the *network-capable DAC* which can become a full-fledged music server with the addition of a network server that provides music storage and can handle the networking functions (Fig.8-4). This system is available with music storage built into the network server (either hard-disk drive or solid-state memory), or with no storage. In the latter case you simply add an NAS drive. This network server can stream music to any device on the network, including DACs in other rooms of the house. The system is controlled by an app, allowing you to play any music from your library in any room with a DAC on the network. The network can be wired or wireless (WiFi).



Fig.8-4 A network-capable DAC that can become a full-fledged music server. (Courtesy Naim Audio)

A DAC that can accept streaming audio from a server via a network is also called a *renderer*. Such a DAC will usually be *Universal Plug and Play* (UPnP) compliant, meaning the DAC adheres to certain established protocols that allow all UPnP-compliant devices on the network to work together seamlessly. With UPnP, no user setup or intervention is required for the devices on the network to communicate.

As you can see, you have many options for accessing file-based music, and not all of them require a computer.

Turnkey Server Considerations

When choosing a server, take into consideration the amount of disk-drive capacity supplied with the machine, along with the cost of adding drives. Some servers allow you to add any relatively inexpensive, off-the-shelf NAS drive; others force you to buy expensive proprietary hardware from the server manufacturer.

Some turnkey servers are designed as “lifestyle” products rather than as serious attempts to deliver high-end sound quality. That is, they exist for customers who want the convenience of playing background music throughout the house. These products tend to have outstanding user interfaces, but often have poor sound quality. Their manufacturers assume that the product’s sound quality is “good enough,” and that their customers will choose a server based on features and interface, not sonic performance. Before buying a turnkey server, consider the manufacturer’s orientation and intentions, and audition the server before you buy. Sometimes such a server’s sound quality can be significantly upgraded by adding an outboard DAC of higher quality, to replace the server’s own DAC. To do this, however, requires that the server offer a digital output jack.

As music servers have gained in popularity, high-end audio companies, which have traditionally been in the business of making amplifiers, disc players, and DACs, are increasingly offering servers of their own. But designing and manufacturing a music server is a daunting task, even for a company that’s highly computer-savvy. Such a company might not have the resources to invest in developing a great user interface, instead coming to market with something of marginal utility. There are examples of traditional high-end companies developing great servers, but be aware that bringing to market a music server requires a completely different skill set from creating, for example, a preamplifier.

A major consideration when choosing any product that includes a music-management app is the app’s ease of use. Some apps are given short shrift in the rush to get the product to market, and quickly become a source of frustration. Others are carefully thought out by designers who understand how the app will be used, then refine the app through extensive beta testing with audiophiles. How easily an app lets you manage your library and select music can spell the difference between satisfaction and annoyance.

Using a Smartphone or Tablet with a Music Server

Today's smartphones and tablets contain integral music-management apps that allow them to function as the core of a file-based music system. Some even incorporate high-resolution storage and playback capabilities.

A smartphone or tablet can operate as part of a music system in one of two ways. The first is to store the music files on the smartphone or tablet device, and play those files from the smartphone or tablet through headphones, or stream the files to the audio system in your home. The second method is to store the music on a hard drive connected to your computer or home network, using the smartphone or tablet purely as a control interface.

Internet Radio and Music Subscription Services

Many consumer-electronics products, including audio/video receivers, smartphones, tablets, Blu-ray Disc players, and HDTV sets, provide integral Webcasting of Internet radio that delivers radio stations from around the world to your computer via the Internet. The audio is coded with a lossy compression/decompression (codec) system such as Advanced Audio Codec (AAC), and streamed to your computer rather than downloaded as a file. Internet radio gives you access to virtually any radio station in the world.

While Internet radio is free, other services that stream music over the Internet are not. The music industry is moving toward a subscription-based model of distribution in which you pay a monthly fee for unlimited access to virtually any piece of recorded music available from that source. You never own a recording of the music, or even store it as a file on your computer or portable device. Rather, you pay for the ability to access the music. Some services offer free music, but you have to listen to advertising. The upside is, obviously, instant access to the world's library of recorded music. The downside is reduced fidelity dictated by the low bitrates offered by these services. A typical bitrate is 64kbps (64,000 bits per second) with AAC encoding, though some offer higher bitrates of 192kbps or even 320kbps. Still, even these higher bitrates are small fractions of CD's bitrate of 1.41Mbs, and consequently offer much lower fidelity than even CD. In addition to these variations in bitrate and sound quality, subscription services vary in how they let you access music. For example, some don't allow you to listen to an artist's latest album in its entirety, and restrict your ability to repeat the song or "rewind" during play. I've avoided describing specifics of the various services because the features, sound qualities, and pricing models are constantly in flux as the industry tries to find a formula that works. In my view, these services are useful for mobile audio and for discovering new music, but not for serious listening through a high-end system. You can, however, purchase high-res downloads of music brought to your attention by a subscription service.

These subscription-based services are available via most consumer-electronics devices, as well as via Android and iOS devices. Examples of these services include Tidal, Beats Music (bought by Apple in mid-2014), Pandora, Rhapsody, and Spotify.

Spotify offers a bitrate of 320kbps, or about a quarter the bitrate of CD. Pandora is even lower, at 192kbps. Obviously, these low bitrates don't offer high

fidelity, but streaming services are great for discovering new music, as sources of background music, or in a desktop system where ultimate fidelity isn't of paramount importance. Even though most compressed streaming services don't offer great sound, don't dismiss them out of hand. As I wrote in [Chapter 3](#), "don't let being an audiophile interfere with your enjoyment of music, anytime, anywhere." The European streaming service WiMP renamed itself Tidal in 2014 and expanded into North America. Tidal offers full CD-resolution 44.1kHz/16-bit streaming using Apple Lossless compression ([Fig.8-5](#)). If the superior sound quality offered by lossless streaming services attracts customers, it may put pressure on services that offer only low-bit-rate compressed files. One day, any music distribution channel that isn't of at least CD resolution may be a nonstarter, and high-res subscription services are on the horizon.

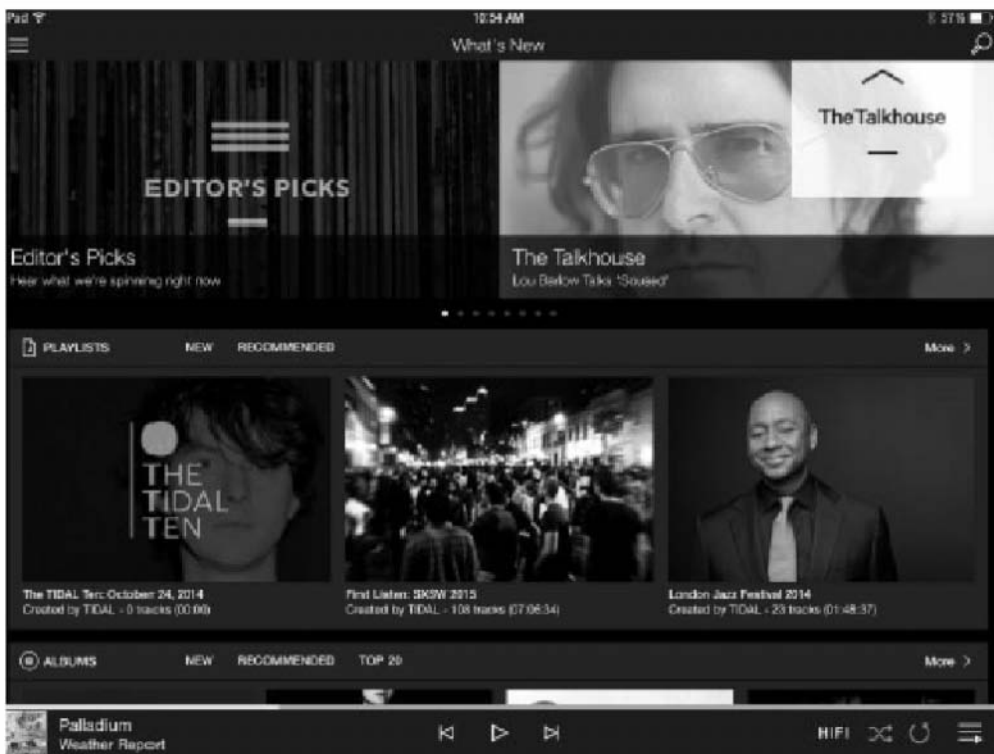


Fig.8-5 The Tidal interface allows you to select music to hear, which is streamed at full CD quality.

Some streaming services offer the ability to store music for later playback, such as when you won't have an Internet connection. The amount of music you may store is, however, restricted, and after being downloaded, the music remains available for only a limited time.

One aspect of some streaming services is their ability to tailor (or curate, in industry parlance) the titles suggested to you based on your musical tastes. With Pandora, for example, you can input the name of an artist or musical genre, and Pandora will deliver music with characteristics similar to those of the artist or genre you selected. In my experience, Pandora is uncannily accurate at doing this, an ability made possible by the Music Genome Project, on which Pandora's software is based. Here's their description of the Music Genome Project, taken from Pandora's Web site:

We believe that each individual has a unique relationship with music —no one else has tastes exactly like yours. So delivering a great radio experience to each and every listener requires an incredibly broad and deep understanding of music. That's why Pandora is based on the Music Genome Project, the most sophisticated taxonomy of musical information ever collected. It represents over eight years of analysis by our trained team of musicologists, and spans everything from this past Tuesday's new releases all the way back to the Renaissance and Classical music.

Each song in the Music Genome Project is analyzed using up to 400 distinct musical characteristics by a trained music analyst. These attributes capture not only the musical identity of a song, but also the many significant qualities that are relevant to understanding the musical preferences of listeners. The typical music analyst working on the Music Genome Project has a four-year degree in music theory, composition, or performance, has passed through a selective screening process and has completed intensive training in the Music Genome's rigorous and precise methodology. To qualify for the work, analysts must have a firm grounding in music theory, including familiarity with a wide range of styles and sounds. All analysis is done on location.

The Music Genome Project's database is built using a methodology that includes the use of precisely defined terminology, a consistent frame of reference, redundant analysis, and ongoing quality control to ensure that data integrity remains reliably high. Pandora does not use machine-listening or other forms of automated data extraction.

The Music Genome Project is updated on a continual basis with the latest releases, emerging artists, and an ever-deepening collection of catalogue titles.

By utilizing the wealth of musicological information stored in the Music Genome Project, Pandora recognizes and responds to each

individual's tastes. The result is a much more personalized radio experience—stations that play music you'll love—and nothing else.

Metadata

Metadata are packages of nonaudio information that accompany an audio file. This information includes the album art, album title, artist names, track titles, track times, and other descriptive information seen on a music-management app. Metadata is a substitute for the album cover, liner notes, and CD booklet of physical formats. The inclusion of metadata is an important aspect of file-based music; without it, you'd be unable to identify the music contained in the file.

When you rip a CD on an Internet-connected computer or music server, the metadata are downloaded from one of several Web sites that serve as a central storehouse, such as AllMusic (formerly All Music Guide), Gracenote, freedb, and Rovi. The amount and type of metadata downloaded and attached to the audio file varies with the site from which the metadata was collected. Some sites offer reviews of the title, extended information about the artist, photos, and much more. In addition to storing metadata with the audio file when you rip the disc or buy a download, some sites provide a wealth of other information that is delivered to you live via the Internet, rather than from the audio file's metadata. This information can be presented to look like extended liner notes with "pages" that you "turn," just as in a CD booklet. The additional information goes far beyond what you'd find in even the most generous liner notes, with up-to-the-minute listings of other titles by the artist, substantive reviews, artist bios, artist photos, background on the musical genre, and other data. There are also listings of similar artists, with live links to information about them and their albums. Moreover, the information doesn't go out of date, as it does with liner notes or stored metadata; it's continually updated, and delivered anew each time you view it.

File Formats

All music-management software for DIY music servers, as well as many of the turnkey systems, offer the choice of storing your music in one of many *file formats*. Your choices of file format and import settings are crucial on several levels; they affect not only sound quality, but also the amount of disk-storage space required for a given amount of music, future compatibility and flexibility, and how metadata are handled.

File formats can be divided into three categories: *uncompressed*, *lossless compression*, and *lossy compression*. The first category, *uncompressed*, consists of two commonly used file formats, waveform audio file format (WAV) and Audio Interchange File Format (AIFF). A WAV file is an *unadulterated* bitstream that is identical to the source data. That is, a WAV file created from ripping a CD contains the identical PCM data that was recorded on the CD. Consequently, WAV files consume the most disk space for a given playing time: 10.5MB per stereo minute for CD-quality audio (44.1kHz/16-bit). The alternative uncompressed file format is AIFF, first developed by Apple Computer. The advantage of WAV and AIFF files, beyond the fact that they sound the best, is that they are “universal currency”: they can be converted to another file format in the future with no loss of quality.

The downside of WAV files is that they have no provision for embedded metadata. Without this metadata, the music-management app will list the title as “Unknown Artist,” and the tracks as simply “Track 1,” “Track 2,” etc. The metadata can be entered manually, but this process is tedious. In addition, the metadata you create manually will be lost if you copy the audio file to another drive, such as for backup. To avoid this limitation, some manufacturers have created proprietary methods of including metadata with WAV files. You could create a large music library as WAV files on an NAS drive with this proprietary metadata scheme, only to discover that the metadata disappear when your NAS drive is connected to any other company’s server. That is, your library is essentially captive to the company’s hardware platform. You can, of course, move the NAS with your WAV files to another server, but the metadata will be unreadable by the new server.

AIFF provides the same advantages as WAV (no sonic degradation, ability to convert to other formats without quality loss), but with permanent metadata tagging. iTunes will convert WAV files to AIFF, but will temporarily double the

hard-drive space needed, until you delete the WAV files. If you've manually entered metadata for the WAV file, that metadata will be attached to the AIFF file. Note that this applies only to standard-resolution files (44.1kHz/16-bit), not to high-res files. iTunes will downconvert high-res WAV files to AIFF, but will reduce the sample rate and degrade fidelity. This is obviously not recommended. If you want to convert WAV to AIFF, use a music player, such as Pure Music, that has "batch mode" file-conversion utilities that will preserve the original sample rate and resolution.

WAV and AIFF are uncompressed formats that consume the most disk space and have the highest bitrates. Let's next look at the category of lossless compression. A codec that losslessly compresses an audio file can almost halve the amount of disk space required compared with uncompressed formats, yet deliver sound with perfect, bit-for-bit accuracy to the original data—which is why such formats are called lossless. Examples of such formats are Apple Lossless Audio Codec (ALAC) and Free Lossless Audio Codec (FLAC). iTunes doesn't support FLAC files, but an aftermarket program such as Pure Music allows FLAC files to seamlessly integrate within iTunes with files of other formats.

You can think of ALAC and FLAC files as similar to Zip files on a PC. Zipping the files produces a smaller file size, but the original files can be recovered later by unZipping them. Although ALAC and FLAC are bit-for-bit accurate to the source data when decompressed, they nonetheless sound slightly inferior to WAV or AIFF files when played back directly from iTunes. How can that be? It is theorized that the process of decompressing the file on the fly during playback requires additional CPU activity, and that this degrades the sound quality. Some music-playing softwares, such as Pure Music and JRiver, decompress the FLAC or ALAC file to PCM, then buffer the PCM data before they're output, resulting in no sonic penalty. As always, there's no substitute for firsthand listening experience when forming opinions about such things. It's easy to compare a WAV and a FLAC file of the same recording and listen for yourself. Incidentally, Windows Media Audio Lossless (WMA Lossless) is another example of a lossless compression format.

The third category of file formats, the lossy compression schemes, are anathema to high-end sound and should be avoided for serious listening. These systems reduce the number of bits in the datastream by throwing away musical information. Unlike uncompressed and lossless-compression formats, which provide perfect, bit-for-bit-accurate reconstruction of the original datastream, lossy compression formats permanently remove information that cannot be recovered. Examples of lossy compression formats are MP3, Ogg Vorbis,

Advanced Audio Codec (AAC), and the Dolby Digital format used to encode film soundtracks. These formats can operate at a variety of bitrates, the higher rates delivering better sound quality. Nonetheless, with the cost of disk storage so low, there's no reason to use a lossy compression format. Note that iTunes delivers music in the lossy compression format AAC at a bitrate of 192 or 320kbps, which is a fraction of the 1.41Mbps bitrate of stereo CD-quality audio.

The default CD-importing setting in iTunes and other players is often one of the lossy compression schemes. It is important that you reset the software's import settings to an uncompressed or lossless file format. You must perform this action every time you download an updated version of the music-player software.

I should mention one more file format: *Digital eXtreme Definition* (DXD). A professional format often used in editing and performing signal processing on DSD files, DXD is simply PCM audio with a sampling frequency of 352.8kHz and a word length of 24 bits. Some consumer products will upsample standard-resolution or high-res audio to DXD. The data rate of DXD is a whopping 8.4672Mbps per channel—more than twelve times CD's bitrate of 705kbps per channel.

The Do-It-Yourself, PC-Based Music Server

A PC-based server can be built on a laptop, notebook, or desktop machine running under the Microsoft Windows or Apple Mac operating systems. Just about any recent-vintage personal computer can be turned into a music server. While a turnkey server can cost five figures, a do-it-yourself system can be had for a few hundred dollars. The disadvantage is that the DIY system requires some technical skill and patience to get the whole thing to work. Some audiophiles, however, see building a music server as an ongoing hobby.

The main components of a DIY server are the computer, a large-capacity hard drive (internal or external), a CD/DVD drive (for ripping CDs), input/output connections (either integral to the computer or via add-on boards), and music-management software. I'll cover the fundamentals of each of these components conceptually. This section is intended as a general guide to DIY music servers, not as specific instructions on how to build one—there are too many variables that will determine any given server's configuration and intended use. Moreover, computer technology advances so rapidly that specific hardware and software recommendations would quickly become outdated.

The Importance of Bit Transparency

You need some degree of computer knowledge to get a DIY server to work and to sound better than a CD source. That's because the hardware and software must be configured perfectly for the server to be *bit transparent*. This term describes the ability of a computer and the music-management software to pass a bitstream from input to output with no change in the data. If the data are changed, the sound will change as well, and never for the better. Data corruption can occur in hardware or software, making it difficult to track down. It's fairly common for DIY servers to degrade sound quality in ways that are hard to troubleshoot. One useful tool for determining whether a server is bit transparent is to connect its digital output to a digital-to-analog converter that is HDCD compatible. Rip an HDCD-encoded CD to the server, play it back, and see if the DAC's HDCD indicator LED illuminates. If it does, the server is bit transparent. The least significant bit in an HDCD-encoded bitstream includes a flag that tells the decoder that this is an encoded bitstream. If the data are corrupted, that flag is lost, the HDCD light doesn't illuminate, and HDCD decoding doesn't happen.

The music-player software sometimes includes obscure settings that you must check off correctly, to avoid having the software corrupt the data. You should also turn off all signal processing in the software. The exception is volume control, which in some cases is better performed in the computer by the music player than in the DAC. Some DACs simply shave off bits to reduce the volume, but the music player can attenuate and redither the signal, reducing the resolution loss. Adding to the confusion, updates of the music-player software can introduce data corruption after you've made sure that the system is bit transparent. Any time you change any hardware or install a newer version of the software, you should recheck that the system is still bit transparent.

Achieving bit transparency is easier with a Mac than with a Windows-based computer because iTunes delivers bit transparency with no user intervention required. Bit transparency on a Windows machine requires that you install a small piece of software called the *Windows Audio Session Application Programming Interface (WASAPI)*. This program bypasses much of the Windows audio processing, so that the music-management software has direct control of the digital audio output. Note that WASAPI is available only with Windows 7 and later operating systems. Another term you'll see is *Audio Stream Input/Output (ASIO)*, another bit of code that bypasses the normal Windows audio processing

so that the music-management program can directly communicate with playback hardware, such as the computer's audio interface.

Computer Requirements

The central processing units (CPUs) of today's computers provide more than enough horsepower to run a music server. But for optimal sound quality, choose a computer platform with the most powerful CPU and the most random-access memory (RAM). The computer should be able to support a large amount of RAM—the more, the better. A large amount of RAM reduces how often the CPU must access the disk drive, making the system run more smoothly.

If the computer is in the same room as the loudspeakers, fan noise is a consideration. The constant background sound produced by computer fans can be a distraction, particularly with classical music recorded in high resolution: The very low noise floor of high-res digital can be ruined by fan noise. The solutions are to put the computer in another room, build a “hush box” around the computer to damp the noise, or buy a fanless PC. The fanless PC has heatsinks on its chassis to draw heat away from internal components. One of my own PC-based servers is fanless. It's also driveless, having 64GB of solid-state memory instead of a hard drive. This makes the server perfectly silent—it has no moving parts. Hard drives produce their own noise, either from the spinning disks or from their own internal fans (in some external drives). My second server is a MacBook Pro with a Synology NAS drive, which has a low-speed fan; I've never found the fan noise intrusive.

Operating System and Playback Software

Your choice of a Mac or PC hardware platform will dictate the operating system (OS) under which your music-player software will run. Generally, it's easier to set up and achieve bit transparency with Mac-based systems running an Apple OS than with PCs running Windows. iTunes running under Windows will resample files of all sampling rates to the rate you've manually selected in iTunes. To achieve bit transparency and avoid sonic degradation, you must manually reconfigure the audio settings every time you play a file with a different sampling rate.

In Windows machines, this awkward step can be circumvented by choosing a soundcard supplied with hardware-specific drivers that enable your music-playback program to bypass the OS's limitations. Note, however, that some music-playback programs, most notably iTunes, lack the essential functionality to access these dedicated hardware drivers. To ensure that you can take full advantage of the capabilities of your audio interface hardware, look for playback programs that support ASIO or kernel streaming, such as MediaMonkey, JRiver Media Center, and foobar2000.

The only currently available workaround for the sample-rate-switching problem in Mac OS is to purchase a software program such as Pure Music, which essentially duplicates the hardware-driver functionality provided free with the USB or FireWire audio interfaces used on the PC platform.

These scenarios are likely to change as technology progresses. When selecting the primary components that comprise a DIY music server, it is imperative to consider how these elements ultimately will work together.

The music-management software organizes your music files, provides the interface for importing and playing music, assembles playlists, converts file formats, and performs other useful functions. The most popular and familiar music-management software is Apple's iTunes, which is free, easy to learn, and reliable. Moreover, if you already use iTunes, you won't have to go through the learning curve required by a different software. The downside is that iTunes is structured around artists and songs rather than composers, orchestras, and conductors, making it less than ideal if your music collection is largely classical. iTunes has another drawback for the serious music enthusiast: It is not designed to handle high-resolution files. It will often downconvert high-res files to standard-res, or corrupt the high-res datastream.

An alternative to iTunes and to Windows Media Audio (which is part of Windows) is third-party music-management software. Popular choices for Windows-based servers are MediaMonkey and foobar2000, both of which are free (retailed versions with upgraded feature sets are also available), and JRiver Media Center.

iTunes on the Mac can be greatly improved in sound quality and functionality with the addition of software such as Pure Music, from Channel D Software. Pure Music uses the iTunes user interface for music-management functions, but all audio processing is performed by Pure Music. The advantages of Pure Music include automatic switching of sample rates (no manual settings required when going from standard- to high-resolution playback, for example), the ability to see and access FLAC and DSD files in your iTunes library (FLAC and DSD are not supported by iTunes running alone), and signal-path optimization for the best sound quality. Pure Music also buffers the audio file during playback for improved sound quality. On a computer running Pure Music, it looks as though iTunes is your playback software and Pure Music is running in the background. In reality, Pure Music is handling all playback functions; iTunes is used solely for its familiar user interface.

The free Remote app, available from the Apple store, provides the iTunes interface for your iPhone or iPad, to allow you to control playback from your listening seat.

Hard-Disk Storage

The DIY server can use the computer's internal disk drive for storing music, or rely on one or more external drives. How much music-storage capacity you need depends on the amount of music you plan to store, the resolution of the files, and the file format selected. What follows assumes that you're not willing to sacrifice sound quality for greater storage capacity, and will store your music as WAV, AIFF, FLAC, or ALAC files. There's no reason to use inferior-sounding formats such as AAC, WMA, or MP3, because storage has become so inexpensive and is getting cheaper all the time.

As described earlier, WAV and AIFF formats store the original PCM datastream with no compression. WAV files consume 10.5MB per stereo minute (630MB per hour) for CD-quality audio (44.1kHz/16-bit). That means, for example, that a 1TB disk drive can store about 1600 hours of music. High-res WAV files encoded at 96kHz/24-bit consume about 35MB per stereo minute (2GB per hour), reducing the disk's capacity to 500 hours. The highest-resolution files further lower this capacity: 192kHz/24-bit files consume 70MB per stereo minute (250 hours of storage on a 1TB disk). In short, the more music you want to store, and the higher the resolution of that music, the greater the storage capacity you'll need. When deciding on disk capacity, allow plenty of space for expanding your music library—it's amazing how fast disk space is consumed. Finally, check to see if the drive's total capacity is provided by a single drive inside the case, or if that capacity is divided among multiple drives. The single-drive option is preferable because there's less chance for drive failure—there's one fewer drive to fail. With a very large drive, one factor you should be aware of is that, compared with a smaller drive, the operating system may take several seconds longer to catalog and display your entire music library if the drive has been in "sleep" mode. This could lead to erratic performance.

When allocating your music-server budget to a hard drive, consider that you should have at least one backup drive, in case the main drive crashes—and, sooner or later, it will. It's a fact of life that hard drives crash and can't be resurrected. Your music library on a hard drive represents a considerable amount of work; protect that investment of time by backing up your library on a redundant external drive. Most drives can be purchased as pairs in a *redundant array of independent disks* (RAID) configuration, in which the second drive is a continuously updated duplicate of the first. Although a RAID is better than a single-disk drive, it's not a true backup: If you inadvertently delete a file, it will

also be deleted on the duplicate array. For true backup, you need a physically separate hard drive or NAS, and use it often. This is a cheap insurance policy that will protect you from having to reload your entire music library if your primary drive crashes.

The disk's rotation speed and amount of onboard buffer memory are also important factors. While faster hard drives make sense for general computer operations, the latest "green" (low-power-consumption) hard drives, which spin at 5400 or 5900rpm, are a better choice for a music server than 7200rpm disks; the data-transfer rate from a 5400rpm disk is plenty fast enough for high-resolution stereo audio, and the slower disk speed allows the drive to run quieter and cooler. Slower drives are also less expensive. Drives with smaller platters (2.5" as opposed to 3.5") tend to be quieter, all else being equal. All disk drives have integral memory that acts as a temporary storage area between the CPU and the drive. A larger data buffer allows the system to operate faster.

A fundamental choice is whether to use a drive inside your computer, or one or more external drives. An external drive is always better, for several reasons. At some point, the internal drive that also holds the operating system will run out of space, and you'll need to transfer the music files to an external drive. It's much simpler to copy files from an external drive to another external drive than from an internal drive to an external drive. That's because copying the contents of an external drive to another external drive preserves the file organization, something that's not true when copying files from an internal drive to an external drive. Another benefit of storing your music library on an external drive is that it can be easily moved to a new computer. Today's drives are fast enough to handle high-resolution digital audio, which requires a minimum data-transfer rate of 80MB per second. This rate is easily achieved by NAS drives with eSATA, Ethernet, FireWire, Thunderbolt, or USB 3 connection.

A related issue is where to store the computer's operating system. Ideally, the computer's OS will reside on a solid-state drive (SSD), with the music library stored on a high-capacity external hard drive. Solid-state memory is lightning fast—my PC-based server and MacBook Pro boot up in a few seconds from their SSD.

Two terms related to hard drives you should know are *network attached storage* (NAS) and *direct attached storage* (DAS). A DAS drive is an off-the-shelf drive connected to your computer via an ATA, SATA, SCSI, USB, or similar interface. The DAS drive is simpler than NAS but lacks the latter's networking capability, its data reading and writing instead being controlled by a single source: the host computer.

An NAS drive is more sophisticated than a standard hard drive you would buy for backing up, for example, your word-processing or photo files. An NAS drive is a computer in its own right (though one of limited capacity) that includes rudimentary software. This additional capability allows the NAS to communicate with other devices connected to the network and share files between them. You can also access NAS drives through a Web browser. NAS drives are often modular, with two or more raw drives housed in a case that also contains the networking software and interface. In an NAS array, one or more of the drive mechanisms can be configured as a redundant drive, improving reliability. Nonetheless, it's essential that you have a physically separate drive for backing up your music library; don't rely only on the NAS array's redundant drive. Fig.8-6 shows a popular NAS drive used for music storage.



Fig.8-6 A dual-bay NAS drive designed for networked media (Courtesy Synology)

NAS drives are better for music servers because of their abilities to be connected through a router to other devices and to share data or storage capacity with those devices. Today's NAS drives are much easier to set up than those of a few years ago, making the question of DAS or NAS moot. NAS drives, however, cost much more than DAS drives—about four times the price for the same storage capacity. You'll need to buy the diskless NAS bay and the drives separately, although some retailers offer kits of bay and drives.

One type of audio-specific NAS, called a *ripping NAS*, includes a CD drive that allows you to rip CDs directly to the NAS. If you opt for a ripping NAS, be

aware that most don't offer the ability to verify that the ripped file is identical to the source file. Moreover, some ripping NASes may introduce bit errors in the ripped file. The best way to ensure that you've achieved a bit-accurate rip is to use a program such as Exact Audio Copy, dBpoweramp (for Windows), or Max (for Mac). Many novice users are unaware that the Use Error Correction feature in iTunes merely activates one rudimentary error-checking feature and doesn't guarantee accurate rips. Ripping software is described in more detail later in this chapter.

If you don't need huge storage capacity, you might consider storing your library on a solid-state drive. One advantage of an SSD is that it has no moving parts and thus offers perfectly silent operation. It's also blazingly fast; your computer will boot up virtually instantly from an SSD, and accessing music is instantaneous. The downsides are limited storage capacity and high cost. However, an SSD used with a fanless computer results in a server with no moving parts whatsoever. Since most SSDs incorporate a standard SATA interface, they can also be installed in inexpensive external drive enclosures that connect to the computer via FireWire, eSATA, or USB.

Loading Your Server with Music

There are generally three ways of getting music into the server: by ripping a CD (transferring music from the CD to the computer's hard drive), buying music downloads from an Internet music site, or paying a service to rip your music library to disk. If you choose to rip your own library, you can use the standard ripping software that comes with your computer's operating system or music-management software, or opt for specialized software such as Exact Audio Copy (EAC), dBpoweramp (for Windows), or Max (for Mac). These programs read data from the CD and can convert the files into a variety of formats suitable for disk storage, as well as into virtually any other format, such as AIFF or FLAC. Their major advantage over conventional ripping software is their ability to ensure verifiably accurate rips. First, they can detect errors in the datastream, rereading error-prone sections of the CD multiple times, until an error-free bitstream is obtained. They can even deliver a perfect bit-stream from damaged CDs, although the process might take several hours. Perhaps even more significant, they can provide you with confirmation of rip accuracy. This is accomplished through several methods. The most common is for the program to rip the music twice (it automatically deletes the first "test" rip when finished) and then compare the two rips; if they are identical, then they are probably an exact representation of the data on the CD. If they are not identical, the program rereads the data from the CD. Another software, called AccurateRip, compares your CD rip to an Internet database of other rips of the same disc to check the accuracy of your rip.

Downloading music from a legitimate Internet music site is a great way of acquiring music and expanding your music library. Be aware, however, that the audio quality can vary considerably. Some sites offer music encoded only with one of the lossy compression formats, resulting in grossly inferior sound quality. Other sites advertise "CD quality" downloads, but you should verify that claim before spending your money. The term "CD quality" has been misused to mean everything from low-bit-rate MP3 files all the way up to true CD quality (44.1kHz/16-bit with no compression). True CD quality is the minimum standard you should accept. Even better are high-resolution downloads that deliver better-than-CD quality for a small increase in price. Some sites offer real 96kHz/24-bit or 192kHz/24-bit downloads sourced from high-resolution master

recordings. When played back on a bit-transparent server and decoded through a high-quality DAC, these files can sound spectacular—and far better than CD.

Some Internet music sites advertise “high-resolution” and “96/24” quality, but the files they supply are nothing more than standard-resolution files that have been upconverted to 96kHz/24-bit. These files will consume more than three times the disk space of 44.1kHz/16-bit files and cost more to buy, but offer no improvement in sound quality over standard-resolution files. Once a digital audio signal has been subjected to 44.1kHz/16-bit encoding, the damage has been done. Choose a download site that has a reputation for high audiophile standards, and whose titles are regularly reviewed in high-end magazines. Some of these sites wouldn’t think of offering anything less than the best sound quality possible.

Finally, you can pay an outside company to rip your music library. You simply ship your CDs and a blank hard drive to the company, and they return your CDs, along with the music on those CDs ripped to the drive. These companies will also rip your CDs to DVDs or to a portable music player. You must take the CDs out of the cases and load them onto spindles supplied by the company, a tedious process. Be sure that the ripping service offers ripping to an uncompressed or lossless compression format. While you might be inclined to choose uncompressed WAV or AIFF files because they are universal standards that can be later converted to any other file format, neither the WAV nor AIFF format incorporates a universally compatible tagging structure for storing the music’s metadata. For this reason, it may be more prudent to choose one of the lossless formats—e.g., FLAC, WMA Lossless, or ALAC—that provide robust, widely supported, and easily converted metadata tags. Remember, however, that these ripping services are based on fast turnover and may not provide documented evidence of the accuracy of their rips, such as that provided by a program such as Exact Audio Copy or dBpoweramp.

An important issue with any music server, whether turnkey or DIY, is the accuracy of the metadata applied to the ripped files. As described earlier in this chapter, metadata are packages of nonaudio data attached to the audio file that identifies the artist, the track, the playing time, and other information. When you rip a CD to iTunes, for example, iTunes automatically downloads this metadata from a service such as Gracenote, AllMedia, Rovi, or freedb. AllMedia is widely regarded as the most accurate tagging service, and also has the highest accuracy and greatest depth of information. These services are generally free to consumers; the companies generate revenue by licensing their services to consumer-electronics manufacturers. When ripping a CD, the service compares the table of contents of the CD being ripped to its database of tens of millions of CD titles.

It's impossible to overstate the importance of accurately tagging your library, and organizing your music in a logical and consistent way. All entries in your library should be tagged using the same tagging conventions as all the other entries, such as the artist's last name, followed by the first name. The various tagging services follow different conventions, resulting in inconsistent organization of your music. Moreover, the tags applied to a file are sometimes incorrect, or contain small errors. You must therefore monitor and manually edit the tags for consistency, and to correct mistakes. The PerfectMeta feature from dBpoweramp gathers metadata from multiple online databases (including AllMedia) and intelligently compares the results to create a tag that is most likely to be correct.

Accurately tagging classical music is a challenge because the industry-standard tagging structure is designed for popular music (artist, album, song) rather than for the title of the work and the names of the composer, conductor, orchestra or ensemble, and soloist(s). Some proprietary music servers have developed tagging structures designed specifically for classical recordings, but the tagging software is specific to the server hardware. You must decide in advance the convention you would like to use, making sure that the convention selected is transportable between different music-playback programs. Alternately, you can use extended nonstandard tags, or use a music-playback program's proprietary database structure to store the additional information. The danger of creating a large library of tags in this way is that if you later switch to a different music-playback program, the new program won't be able to read the tags.

Getting Music from the Server to Your Playback System

Now that we've covered getting music into the server, let's look at options for getting it out to your hi-fi system for listening. The simplest method is to connect a USB cable between the computer's USB port and a USB-capable DAC. (USB is covered in detail in the previous chapter.) The USB DAC's analog outputs feed one of your preamplifier inputs. If your DAC lacks a USB input, you can buy a USB-to-S/PDIF format converter. This small, inexpensive device takes in USB from the computer and converts it to S/PDIF on an RCA jack. For state-of-the-art computer audio, a high-quality USB-to-S/PDIF converter has some sonic advantages, including isolating the DAC from the computer's noise as well as reclocking the signal. The downside of bypassing your DAC's USB input with an outboard USB-to-S/PDIF converter is that the DAC is no longer the master clock to which the source must lock. Instead, the clock must be recovered from the S/PDIF, a process that introduces jitter. Whether a DAC sounds better driven by USB or by a USB-to-S/PDIF converter is determined by the quality of the USB implementation and the quality of the USB-to-S/PDIF converter. USB implementations vary in sound quality (although less so than a few years ago), even among otherwise sophisticated designs.

Some computers offer S/PDIF output on a TosLink optical jack, but this interface degrades sound quality.

Desktop computers can be fitted with a soundcard that outputs S/PDIF on an RCA jack or AES/EBU on an XLR jack. These soundcards replace the digital audio output on the computer, which is usually USB but sometimes S/PDIF coaxial or TosLink optical. Replacing the computer's stock digital output with a specialized soundcard usually results in better sound because the specialized card has better output circuitry, more precise clocking, and lower jitter than the computer's own audio output.

Other soundcards include digital-to-analog converters that output analog audio on RCA jacks. These cards don't require an outboard DAC, but are generally limited in sound quality. The inside of a PC is a hostile environment for analog signals; noise can contaminate the analog output. Moreover, the soundcard offers the designer very little circuit-board space with which to work, which limits power-supply sophistication. Although some soundcards with

integral DACs are designed for music playback rather than just gaming or general-purpose computing, they do not provide the ultimate in sonic performance.

The third option is to stream music wirelessly from your computer to your DAC. You'll need WiFi capabilities in your computer, along with a WiFi base station to convert the WiFi signal to S/PDIF (either TosLink optical or coaxial). WiFi is subject to signal dropouts, is limited in distance (25' is generally considered the maximum range), and is affected by your home's construction.

Some DACs offer built-in WiFi, allowing your computer to stream audio to the DAC without the need for any physical connection. A huge advantage of this approach is decoupling the noisy computer from the sensitive DAC circuitry. In an electrical connection such as USB, noise on the computer's ground travels down the USB cable to the DAC's ground, degrading sound quality. In my own comparisons of wired (USB) and wireless DAC connection, wireless sounds significantly less grainy, opaque, and harsh. Nonetheless, the WiFi signal must be robust to avoid dropouts, or the computer and DAC losing the connection between them which must then be reset manually.

Networking

Your music server can function as a standalone device or be connected to a home network. If you use an NAS drive for music storage, you've created a network, although one that may be limited to the computer and the NAS drive. The decision is whether to connect this network to other devices in your home so that they can share content between them. For example, a home network allows you to access your music library stored on the NAS from a network-compatible audio system in a bedroom. Any device on the network can communicate and share files with any other device on the network. This network can be wired or wireless, but a wired connection is more secure and reliable.

Music files are served to the network by devices on the network. A source device, called the *server*, streams files to *client* devices on the network. (The term *server* has a more specific meaning in the context of a network system than the general term *music server*.) The example above of playing (streaming) music files stored on your main system's computer (the server) on a bedroom system (the client) is typical. Some client devices don't require a computer to stream the audio; the client pulls files from an NAS drive on the network and decodes them for listening.

It isn't necessary for your network to be connected to the Internet, but it's a good idea—an Internet connection is required to gather metadata when ripping a CD. Obviously, you'll also need an Internet connection for downloading music files and updating software.

Two networking terms you should know are *Universal Plug and Play (UPnP)* and *Digital Living Network Alliance (DLNA)*. UPnP is a set of networking protocols that allow connected UPnP devices to recognize each other and to share data between them. DLNA is built on the UPnP platform, but is more restrictive regarding the content that can be shared. DLNA does this to streamline the technology and make it more reliable for consumers, and offers more features than UPnP.

DLNA is a trade group that has developed protocols for sharing multimedia among devices on a network. When you connect a DLNA-certified component to your home network, it automatically connects to and communicates with all other DLNA-compliant devices on that network, regardless of the components' brands. I've long lamented the sad state of affairs in consumer electronics in which audio and home-theater systems are built from "dumb" boxes that operate

in isolation with no knowledge of other components in the system, and with each controlled by its own remote control. DLNA solves this problem by making various devices, made by disparate manufacturers, work together seamlessly. For example, you can stream music over your network from a DLNA-certified NAS, smartphone, or tablet to a DLNA-certified desktop audio system in your kitchen. Similarly, you can use a smartphone to remotely control streaming from a PC to a bedroom system. The possibilities are virtually endless.

A Simple, Do-It-Yourself Music Server

I recently created a Mac-based server that was easy to build, configure, and use. Plus, it sounds great at the front end of a reference-quality system.

I began with a MacBook Pro, but you could base your system on a Mac mini or any other Mac. Ideally, the Mac would be a dedicated music server that would let you optimize its performance for music. If you're buying a new computer, opt for one with the fastest processor and the greatest amount of RAM you can afford, and specify a solid-state rather than a hard-disk drive. This SSD will store the computer's operating system, along with iTunes and your music-management software. The advantage of an SSD is its silent, near-instantaneous operation.

There's no reason to spend money on a big screen with high resolution—a 13" model will work just fine. If the computer doesn't offer an integral optical drive, you'll need to add an outboard optical drive for ripping CDs. When ripping many CDs, loading and unloading CDs in an external drive is easier than inserting them in the computer's integral drive. A third-party ripping program such as Max produces better rips than the iTunes ripper.

You'll store your music library on an NAS drive connected to the Mac via an Ethernet cable. NAS drives can be purchased as kits that include the NAS server, plus two or more raw drives that you'll mount in the NAS server. Many NAS drives are designed for massive storage and multiple users, but you won't need anything that elaborate—a two-bay NAS loaded with a pair of 2TB drives should be sufficient. You'll actually have just under 2TB of storage, because the second drive will merely duplicate the first drive. This RAID configuration, described earlier, shouldn't be relied on for backup. For true backup, you'll still need another physical drive. That drive needn't be an NAS; an inexpensive USB drive is adequate.

Some NAS servers offer integral RAM for faster operation. Another NAS feature is compliance with DLNA. If you connect your server to a network, you may need to add a small, inexpensive device, called a gigabit switch, that expands the number of Ethernet ports on the network. The gigabit switch is a hub, with an Ethernet input and four or more Ethernet connection points.

Once the hardware is connected, buy a license to Channel D Software's Pure Music and download the program. As described earlier, Pure Music replaces iTunes as the playback engine, for better sound and increased functionality.

You'll still use iTunes as the user interface, but Pure Music handles all the playback functions.

Next, download the Remote app from the Apple store. This app allows you to see your music library on your iPhone or iPad, select music, and control playback.

In iTunes Preferences, select the NAS drive as the storage destination for ripped files. If you already have an iTunes library stored on a Mac or PC that is also on the network, you can access that library via iTunes using Apple's Home Sharing app. It's better, however, to copy the music files from the PC or Mac to the NAS. The files on the NAS can be more quickly accessed.

For optimal audio performance, follow these configuration guidelines, each of which will improve your server's sound quality.

- Rip CDs to the AIFF file format, using a third-party program such as XLD or Max. WAV and AIFF offer identical sound quality, but AIFF is a better choice for Macs because the format can store the album information and cover art as metadata along with the audio file. If you rip CDs using the iTunes ripper, be sure to enable Error Correction, choose Automatic, and select AIFF as the file format.
- Go to About This Mac>System Details>USB to see which device is connected to each USB port. Connect your DAC to the USB port with the least number of devices (such as keyboards and trackpads) connected.
- Be sure that no programs other than your music-playback software are running.
- Some listeners report a sonic advantage to turning off WiFi and Bluetooth. Note, however, that the Remote app, which allows you to control playback from an iPad or Apple mobile device, requires that WiFi be enabled on the Mac.
- Go to System Preferences and disable: search functions (Spotlight), Sleep Mode (Energy Saver), and sharing (Sharing).
- Go to the Audio Midi Panel and verify that the output level is set to maximum.
- In Pure Music, engage the Cache/Memory playback mode.
- In Pure Music, check Disable DSP.

- Use high-quality Ethernet and USB cables, and use vibration-isolation devices under the computer and NAS.

Transferring an LP Collection to a Server

Once you have a music server up and running, it's inevitable that you'll want to digitize your LP collection (if you have one) so that you can access that music via the server. Although the prospect of subjecting analog sources to digital conversion is anathema to many audiophiles, it's nonetheless handy to have your entire music collection—from analog *and* digital sources—on the server. I recommend, however, that you keep your LP collection and turntable so that you can continue to enjoy the sound of analog for those focused listening sessions.

Transferring your LPs to digital can be a fairly quick and simple process, or an involved and exacting procedure. As with everything else, how much effort you put into the transcription process will determine the quality of the result. Big-box retailers sell “USB turntables” for \$99 that plug into your computer; your only contribution is a few mouse clicks. At the other end of the spectrum is an elaborate process of cleaning the record, fine-tuning the turntable and cartridge, buying a high-quality *analog-to-digital converter* (ADC), installing high-end LP-transcription software, and recording your LPs in high resolution. We'll cover the various methods and technologies, but keep in mind that you'll probably transfer your LP collection only once. A listener's LP library represents a lifetime of musical exploration; preserving as much of that legacy as possible should be a high priority.

I can't recommend highly enough the step of cleaning your LPs before converting them to digital files. If you don't have a record-cleaning machine, this is a good excuse to buy one. Not only will your LPs sound better for normal listening, they'll produce much cleaner-sounding digital versions for your server. [Chapter 9](#) covers record-cleaning machines in detail.

The quality of the turntable, tonearm, cartridge, and phono preamplifier will be a significant factor in the final sound quality. If you have an inexpensive LP front-end and many treasured records, consider renting or borrowing a better 'table. Fine-tune the *vertical tracking force* (VTF) and *vertical tracking angle* (VTA) before beginning the transfer, and adjust the platter's speed with a strobe (as explained in [Chapter 9](#)). Phono cartridges have widely varying frequency responses, transient fidelity, soundstaging, and other characteristics—remember that these qualities will be imprinted on every LP you transfer. Similarly, the phono preamp's sonic signature will be overlaid on your library of digitized LPs. Be sure to optimally set the phono preamp's *cartridge loading*. If it's a tubed

model, consider replacing the tubes with specially selected low-noise versions before a transcription session. The phono preamp can be a standalone model or part of a full-function preamplifier or integrated amplifier.

Some PC audio interfaces allow you to connect the output from your turntable directly to the computer. This audio interface, once called a soundcard, can be inside your computer or be a small external device. The audio interface amplifies the very low-level signal from the cartridge, digitizes the signal, and performs *RIAA equalization* (explained in [Chapters 4 and 9](#)) via the software supplied with the audio interface. This obviates the need for a phono preamplifier and ADC, but is undoubtedly of lower quality than an audiophile-grade phono stage and ADC. If you opt to use a computer's audio interface, be sure that the interface has enough gain for your cartridge's output voltage. If the interface has insufficient gain, the resulting digital files will be noisy and lacking in dynamic range.

If you choose the higher-quality path of a separate phono stage and ADC, you'll connect the phono stage's line-level output to the ADC's input. Note that analog-to-digital conversion can take place in a standalone ADC, the ADC built into a digital recorder, or the ADC in a computer's audio interface. A few phono stages include an integral ADC, simplifying the process of turning LPs into digital files.

Before choosing an ADC, you should decide on the resolution at which you want to encode your LPs. The lowest resolution you should consider is a 44.1kHz sampling frequency with a 16-bit word length. This is true CD-quality resolution. But with the ready availability of high-resolution ADCs and cheap storage, there's no reason not to digitize your LPs with the best possible quality. Keep in mind, however, that files made with 96kHz/24-bit resolution consume more than three times as much disk space as do 44.1kHz/16-bit files. Be sure the ADC you choose can support the maximum resolution you want.

An important consideration is the ADC's ability to adjust and monitor the signal level. If the signal level feeding the ADC is too high, musical peaks will sound distorted. If the signal level feeding the ADC is too low, you'll be throwing away resolution. For example, if the highest signal level is 12dB below the maximum allowable level without distortion, you're losing two bits of information. If you began with 16 bits, it's like having only a 14-bit converter. Setting the level so that the highest peak on an LP uses nearly all the bits without overload is vital to making a great-sounding transcription.

The ultimate way to rip an LP library is to buy a dedicated LP-transcription phonostage with a flat output—that is, one that doesn't perform RIAA equalization on the analog signal. After this flat signal is digitized, RIAA equalization is realized in the digital domain by software. This technique has several advantages, primarily an increase in resolution of 1–2 bits at lower frequencies. This additional resolution is achieved because the high-frequency filtering of RIAA results in a longer word length. The increase in resolution with high-frequency filtering is identical to what occurs in converting DSD to PCM: DSD's 1-bit code at 2.8224MHz is converted to a 24-bit code at a lower sampling frequency, such as 88.2kHz. Digital-domain RIAA equalization also eases the requirements for perfectly setting the ADC input levels, as described in the previous paragraph; the additional 6–12dB of headroom afforded by the extra resolution makes the level setting less critical. Finally, digital-domain RIAA subjects the audio signal to less analog circuitry, bypassing the capacitors and resistors in a passive RIAA implantation, and even an additional gain stage in an active RIAA circuit.

You'll need software to manage the transfer from the ADC to the computer's disk drive. Again, be sure that the software is compatible with your computer and operating system, and will support high sample rates if you choose. Virtually all software packages include equalization, hiss filtering, pop removal, and other signal processing. Although some of these can improve the sound of the transfer, these "improvements" often come at the price of reduced sonic transparency.

After you've transferred an LP side to a digital file, you'll need to divide that file into individual tracks and label them. Pure Vinyl provides a patented graphic interface that looks like an LP side, complete with visual bands between songs, that allows you to easily find the song starting points (Fig.8-7). Other transfer-software packages recognize the gaps between songs and automatically separate the songs into separate files.

Undoubtedly the most tedious—but essential—step in converting LPs to digital files is tagging those files with the names and titles of the artists, albums, and songs. The same problems noted earlier regarding the accuracy of metadata tags, and whether the library is organized in a logical and consistent manner, apply to tagging files transferred from LP. When creating your own digital music files from LP, you can create this metadata yourself by typing it out on a keyboard, or you can buy an additional piece of software that allows Gracernote (described earlier in this chapter) to identify the music and attach the tags. After the file has been tagged, your file transferred from LP will appear as just another

title on your music server—one ready to play with a few mouse clicks, or some finger taps on a touchscreen.

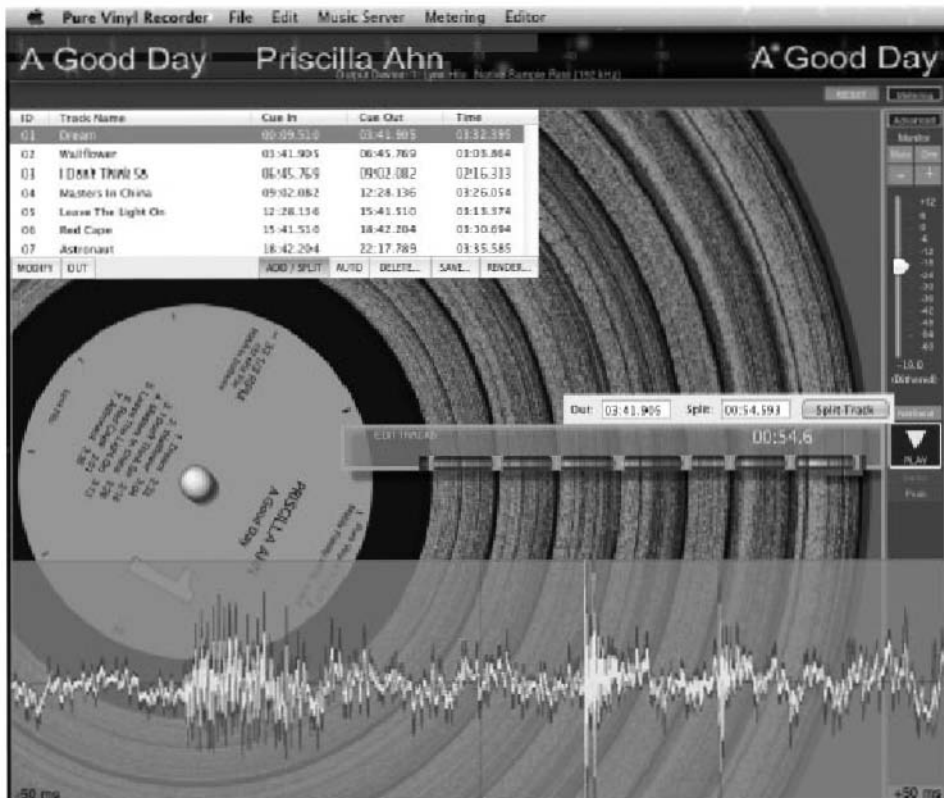


Fig.8-7 Ripped LPs can be separated into individual song files with a graphic interface. (Courtesy Channel D Software/Pure Vinyl)

9 Turntables, Tonearms, and Cartridges: The LP Playback System

Introduction

Beginning in 2007, a curious phenomenon emerged: after years of precipitous decline, sales of LP records began to slowly increase. The LP was considered a defunct format, and even became a synonym for *any* antiquated technology. But that barely perceptible rise in LP sales in 2007 over the previous year was just the beginning of what would become a landslide of monumental proportions. Sales of vinyl records, turntables, cartridges, phonostages, and all things related to playing LPs skyrocketed beyond even vinyl enthusiasts' most optimistic projections. Although records still account for only a small percentage of all music sold, the fact that this archaic technology is thriving in a world of digitally sourced music is remarkable.

The resurging popularity of the vinyl record is a testament to its fundamental sound quality. Not only do LPs deliver a musical realism unmatched by standard-resolution digital audio, they also represent a completely different approach to music listening. At one extreme is the person who downloads compressed files at the lowest possible bit rate, and listens to them in the background on a cheap portable player and \$5 earbuds while engaging in some other primary activity. Listening to an LP couldn't be more different. Placing an LP on a turntable signifies your intention to set aside all other activities and turn your full attention to the music.

The high-end audio industry has responded to vinyl's spectacular comeback with a flood of interesting new turntables, tonearms, cartridges, and phono preamplifiers, not to mention accessories for getting the most out of a record collection. Specialty audiophile labels have licensed classic recordings of the past

and reissued them in magnificent sound. Today's LPs are mastered and manufactured with a level of care unimaginable 20 years ago. The LP was once a cheap commodity; now, making them is an art. But even from the early days of the LP, more information could be encoded in the grooves than was possible to retrieve. Today's LP-playback gear is extracting more and more of that information. The result is that we're now hearing, for the first time, the full measure of the LP format's greatness.

The renewed interest in vinyl playback is largely driven by two demographic groups. The first is young adults who value the LP's sound quality, tactile nature (including the large-format cover art and liner notes), and "retro-cool" factor—qualities that compressed downloads lack. The second group driving vinyl sales is the music lover who grew up listening to records but abandoned them in favor of CDs. After 25 years of listening to digital, these listeners are now dusting off their records and discovering anew the musical pleasures and sonic quality possible with today's LP playback systems.

If you haven't heard a good LP playback system, I encourage you to visit a specialty retailer who can demonstrate the potential of vinyl, or to attend a consumer-audio show where turntables are prevalent. At most audio shows today, the majority of rooms are spinning vinyl.

The general public's perception that CD is vastly superior to LP (remember the "Perfect Sound Forever" marketing campaign for CD?) is perpetuated because very few listeners have heard high-quality LP playback. When done right, LP playback has an openness, transparency, dynamic expression, and musicality that cannot be matched by CD. There's just a fundamental musical rightness to a pure analog source—i.e., one that has never been digitized—that seems to better convey musical expression. There's a saying in high-end audio circles: CDs sound great—if you never listen to vinyl. In other words, until you hear good LP playback, you don't know what you're missing.

This isn't to say that LP playback is perfect. It suffers from a variety of distortions, such as mistracking, ticks and pops, speed instability, surface noise, cartridge frequency-response variations, inner-groove distortion, and LPs' susceptibility to wear and damage. But for many listeners, these problems are less musically objectionable than the distortions imposed by digitally encoding and decoding an audio signal. Moreover, the best turntables, tonearms, and cartridges substantially reduce these annoyances while maximizing the strengths of LP playback. In fact, today's best LP playback systems are remarkable in their ability to make you forget that you're listening to anything but music.

Converting the binary ones and zeros of digital audio back into music has been likened to trying to turn hamburger back into steak. Some listeners can hear past the LP's flaws and enjoy the medium's overall musicality. Other listeners can't stand the ritual of handling and cleaning records—not to mention keeping the turntable properly adjusted—and think CDs are just fine. I think of it this way: LP's distortions are periodically apparent, but separate from the music; digital's distortions are woven into the music's fabric. Consequently, analog's distortions are easier to ignore. If you're inclined to think that the Compact Disc is without fault but have never heard a properly played LP, give yourself a treat and listen to what vinyl can do before you write off the possibility of owning a high-quality turntable.

LP Playback Hardware Overview

The long-playing (LP) record playback system is a combination of a *turntable*, *tonearm*, and *phono cartridge* that converts the mechanical information encoded on vinyl records into an electrical signal that can be amplified by the rest of your playback system. The turntable spins the record, the tonearm holds the cartridge in place, and the cartridge converts the wiggles in the record's groove into an electrical signal. These elements, and how they interact with each other, play pivotal roles in getting good sound from your system.

System Hierarchy: Why the LP Front End Is So Important

My suggestion later in this chapter—that some audiophiles should spend as much as 40% of their total system budget on an LP front end—may strike some as excessive. After all, how much of an effect can the cartridge, tonearm, and, especially, the turntable have on the reproduced sound, compared to the contributions of amplifiers and loudspeakers?

The answer is that the components at the front of the playback chain should be of a quality at least as high as those at the end of the chain (the loudspeakers). The LP playback system's job is to extract as much musical information from the LP grooves as possible, and with the greatest fidelity to the signal cut into the LP. Any musical information not recovered at the front end of the playback system cannot be restored later in the chain. It doesn't do any good to have superlative electronics and loudspeakers if you're feeding them a poor-quality, low-resolution signal from an inadequate LP front end. If the music isn't there at the start, it won't be there at the finish.

The importance of the turntable was first brought to the world's attention in the early 1970s, by Ivor Tiefenbrun, of Scottish company Linn Products Limited. Tiefenbrun sold his turntable by walking into audio dealerships unannounced with the turntable under his arm, and asking the store owner to listen with him to the best system in the store with the store's turntable. After that, they listened to the store's *least* expensive system hooked up to the Linn LP12 turntable. When the store owners heard their most modest system outperform their most expensive system, many became convinced of the turntable's importance. (When asked by store owners, "But how can I sell this to my customers?," Tiefenbrun responded, "The same way I just sold it to *you*.") It took years of these kinds of demonstrations to convince the audio community that the turntable was a significant variable in a system's sound quality. Today, there is no debate about the turntable's importance in achieving good sound. (To date, more than 150,000 Linn LP12s have been sold worldwide.)

How to Choose an LP Playback System

Because there are many more variables to account for in LP playback, how much of your audio budget you should spend on this part of your system is a more complicated decision than setting a budget for, say, a power amplifier. Let's look at two hypothetical audiophiles, one of whom should spend much more on a turntable, tonearm, and cartridge than the other.

Our first audiophile has a huge record collection that represents a lifetime of collecting music. Her record collection is a treasure trove of intimately known music that she plays daily. She has very few CDs, buying them only when her favorite music isn't available on vinyl. She much prefers the sound of LPs, and doesn't mind the greater effort required to play them: cleaning her stylus and records, turning over the record, lack of random access.

The second audiophile's LPs represent a small percentage of his music collection—most of his favorite music is on CD. His LP listening time is a fraction of the time he spends listening to CDs. He likes the convenience of CDs, and can happily live with the sound of his excellent disc player and music server.

The first audiophile should commit a significant portion of her overall system budget to a top-notch LP front end—perhaps 40%. The second will want to spend much less—say, 10 to 15%—and put the savings into the components he spends more time listening to.

In the third edition of this book (2004), I recommended that readers invest immediately in a high-quality LP playback system. I assumed then that, in the future, fewer and fewer manufacturers would offer turntables, tonearms, and cartridges. How wrong I was. Since that edition, we've seen explosions of new LP playback gear at all price levels. Moreover, today's products deliver much better sound at a given price point because of economies of scale in manufacturing and increased competition for the turntable buyer's dollar.

A decent entry-level turntable with integral tonearm runs about \$300, with an appropriate cartridge adding \$60 to \$150 to the price. A mid-level turntable and arm cost \$800 to \$1500. Turntables in this price range include an integral tonearm, and many are supplied with a cartridge. These complete packages are ideal for the first-time turntable buyer; they take the guesswork out of matching the component parts, as well as the work involved in mounting the cartridge and aligning it in the tonearm. [Fig.9-1](#) shows a “plug and play” entry-level turntable.

There are roughly three levels of quality and price above \$2000. The first is occupied by a wide selection of turntables and arms costing between \$3000 and \$6000. In this range, you can achieve outstanding performance. Plan to spend at least \$1000 for a phono cartridge appropriate for these turntables. A turntable/arm combination costing from \$10,000 to \$20,000 will buy you very nearly state-of-the-art performance (if you choose the right model). The sound quality of such a system (Fig.9-2) will be significantly better than that of a turntable-tonearm combination in the lower tiers, and will come surprisingly close to the sound of ultra-exotic models that sometimes rely on bling for their appeal. Later in this chapter we'll consider the sonic differences among turntables, tonearms, and cartridges.



Fig.9-1 Today's entry-level turntables are supplied with a tonearm and cartridge, and come calibrated out of the box. (Courtesy Pro-Ject and Sumiko)



Fig.9-2 A high-quality turntable and tonearm. (Courtesy Basis Audio)

The very highest price level is established by turntables costing between \$25,000 and \$150,000. A topflight phono cartridge can add as much as \$15,000 to the price. These systems are characterized by extraordinary and elaborate construction techniques, such as vacuum hold-down systems to keep the record in intimate contact with the platter, massive vibration-resistant bases, exceptional systems for isolating the platter and arm from vibrations, lots of high-precision machining, and a high-gloss finish (Fig.9-3). Turntables at this price level often provide the option of using more than one tonearm. The second tonearm is often fitted with a cartridge of different sonic characteristics that is best suited to certain kinds of music—or the second arm can hold a monophonic cartridge, for a listener who wants to hear mono recordings without compromise. (Although you can play a mono record with a stereo cartridge, the best sound quality is achieved with a mono cartridge.) Many of these megabuck turntables are sonically superlative, and some turntables in the \$30,000–\$50,000 range are equal (or even superior) to the costliest units in quality of construction, elegance of finish, reliability, and sound quality.

Once you've decided on an LP front-end budget, allocate percentages of that budget to the turntable, tonearm, and cartridge. As I mentioned, most lower-priced turntables come with an arm and cartridge already fitted and included in the price. At the other end of the scale, the megabuck turntables also include a tonearm—turntable and tonearm are engineered as parts of a single system. In the middle range, you should expect to spend about 50% of your budget on the turntable, 25% on the tonearm, and 25% on the cartridge. These aren't hard figures, but approximations of how your LP front-end budget should be allocated. As usual, a local audio retailer with whom you've established a relationship is your best source of advice on assembling the best LP playback system for your budget.



Fig.9-3 An upper-end turntable with a tangential-tracking air-bearing tonearm. (Courtesy Walker Audio)

An item of utmost importance in achieving good sound from your records is a good turntable stand or equipment rack, particularly if the turntable has a less-than-adequate suspension. (Turntable suspensions are described later in this chapter.) I cannot overstate how vital a solid, vibration-resistant stand is in getting the most from your analog front end. Save some of your budget for a solid

equipment rack or you're wasting your money on a good turntable, arm, and cartridge. Note, however, that the very best turntables have extraordinary mechanical systems for isolating the turntable from vibration, and thus don't benefit as much from a solid stand.

The importance of a good rack was underscored for me while I was reviewing a turntable-isolation system. To test the effectiveness of this system, I moved the turntable (a model that lacked a suspension) from its usual place on a 350-pound, sand-filled, spiked turntable stand (the Merrill Stable Table) to a generic wooden stereo rack. The difference was like night and day—and far greater than I'd expected. On the wooden rack, the bass was mushy and colored, transients were less clearly defined, instrumental timbres were synthetic, and there was a general impression of hearing a collection of sounds rather than musical expression.

I'll talk more about turntable stands later in this chapter, under "LP System Setup," and also in [Chapter 15](#). I'll also go into some detail here about how turntables, tonearms, and cartridges are built, and what to look for when shopping for them. You can use this information to assemble your own turntable system, or just brush up on the subject so that you can talk confidently with the dealer who puts together your LP playback system. Whichever route you choose, you'll still want to evaluate the turntable system's sonic merits by listening to it—at length.

What to Listen For

Judging the sonic and musical performance of an LP playback system is more difficult than evaluating any other component. If you want to audition a phono cartridge, for example, you can't do so without also hearing the turntable and tonearm, as well as how the cartridge interacts with the rest of the LP playback system. The same situation applies to each of the elements that make up an analog front end; you can never hear any them in isolation. Although this is true for any component in a system, the parts of an LP front end interact with each other to a greater extent than, say, a preamplifier and a power amplifier. Further, how those three components are set up greatly affects the overall performance. Other variables include the turntable stand, where it's located in the room, the phono preamp, and the load the preamp presents to the cartridge. Nonetheless, each turntable, tonearm, and cartridge has its own sonic signature. The better products have less of a sonic signature than lower-quality ones; that is, they more closely approach sonic neutrality.

A high-quality LP front end is characterized by a lack of *rumble* (low-frequency noise), greatly reduced record-surface noise, and the impression that the music emerges from a silent background. Low-quality LP front ends tend to add a layer of sonic grunge below the music that imposes a grayish opacity on the sound. When you switch to a high-quality front end, it's like washing a grimy film off of a picture window—the view suddenly becomes more transparent, vivid, and alive.

Even if the LP front end has no obvious audible rumble, it can still be adding this layer of low-level noise to the music. The noise establishes a threshold below which musical information is obscured. A better turntable and tonearm strip away this film and let you hear much deeper into the music. It can be difficult to identify this layer of noise unless you've heard the same music reproduced by a top-notch front end. Once you've heard a good LP playback system, however, the difference is startling. To use a visual analogy, hearing your records on a good analog front end for the first time is like looking at the stars on a cloudless, moonless night in the country after having lived all your life in the city. A wealth of subtle musical detail is revealed, and with it comes a much greater involvement in the music.

Another important characteristic you'll hear on a superlative system is the impression that the music is made up of individual instruments existing in space.

Each instrument will occupy a specific point in space, and be surrounded by a halo of air that keeps it separate from the other instruments. The music sounds as if it is made up of individual elements rather than sounding homogenized, blurred, congested, and confused. There are a special realness, life, and immediacy to records played on a high-quality turntable system (if, of course, the recording engineers have captured these qualities in the first place).

A related aspect is soundstage transparency—the impression that the musical presentation is crystal clear rather than slightly opaque. A transparent soundstage lets you hear deep into the concert hall, with instruments toward the rear maintaining their clarity and immediacy, yet still sounding far back on the stage. The ability to “see” deep into the soundstage provides a feeling of a vast expanse of space before you in which the instruments exist. Reverberation decay hangs in space longer, further conveying the impression of space and depth.

Conversely, a lower-quality LP front end will tend to obscure sounds emanating from the rear of the stage, making them seem undifferentiated and lacking in life. The presentation is clouded by an opaque haze that dulls the music’s sense of immediacy, prevents you from hearing low-level detail, and tends to shrink your sense of the hall’s size. The differences between superb and mediocre LP reproduction are much like those described in [Chapter 1’s Grand Canyon analogy on page 2](#).

Other important musical qualities greatly affected by the LP front end are dynamic contrast and transient speed. A top-notch LP replay system has much wider range of dynamic expression than a mediocre one; the differences between loud and soft are greater and more varied. In addition to having wider dynamic range, musical transients have an increased sense of suddenness, zip, and attack. The attack of an acoustic guitar string, for example, is quick, sharp, and vivid. Many mediocre turntables and tonearms slow these transient signals, making them sound synthetic and lifeless. A good LP front end will also have a sonic coherence that makes the transients of instruments playing different notes sound as if they are all lined up in time with each other. The result is more powerful rhythmic expression.

Better LP front ends are also characterized by deeper bass extension and a more solid bottom end. With a great turntable, arm, and cartridge, the bass takes on a whole new level of impact, power, and dynamic agility. Moreover, it’s easier to hear subtleties of bass pitch and dynamics with a good turntable system. Finally, a great LP playback system will reveal a greater richness and density of tone color with instruments such as acoustic bass, electric bass, cello, bass clarinet, piano, and organ.

Just as we want the LP front end to portray the steep attack of a note, we want the note to decay with the actual decay characteristics of the original musical event. When some instruments are played a certain way, there will be a sudden end to the note, such as the sound of a percussion instrument. In other cases, such as a struck cymbal, a piano played without using the damper pedal, or any bowed or plucked string instrument, the decay might last for several seconds. In some cases, the decay will still be very audible even after the next notes played on the same instrument have been struck. This might not seem like a significant factor in the enjoyment of music, but correctly reproducing these decays results in denser, richer tone colors, as the decay of one note changes the character of the next. The composer and/or performer intended that the music be heard this way, and not with a premature truncation of decays. The ability of the best LP front ends to portray and unravel this amazingly delicate and intricate low-level information can be astonishing. Rather than blurring notes together, a first-rate LP front end is characterized by its ability to clearly articulate each note with a sense of silence between notes, while retaining decay information within the inter-note silence. A good test of an LP front end's transient characteristics is intricate percussion music. Any blurring of the music's dynamic structure—attacks and decays—will be immediately obvious as a smearing of the sound, a lack of immediacy, and the impression that you're hearing replicas of the instruments rather than the instruments themselves. Periodically hearing live, unamplified musical instruments can really sharpen your hearing acuity for judging reproduced sound.

All turntables, tonearms, and cartridges influence the sound's overall perspective and tonal balance. Even high-quality components can have distinctive sonic signatures (though the very best analog components can be remarkably transparent). Therefore, to achieve the most musical result, the turntable, arm, and cartridge need to be carefully matched to each other. Matching a bright, forward cartridge to an arm with the same characteristics, and mounting both on a somewhat aggressive-sounding turntable, could be a recipe for unmusical sound. Those same individual components may, however, be eminently musical in a mix of components whose sonic signatures tend to complement each other.

Because, by definition, the very best turntables, tonearms, and cartridges are more sonically neutral than lesser products, system matching is less critical as the quality increases. State-of-the-art turntables and tonearms tend to be so neutral that you can put together nearly any combination and get superlative sound. Even upper-end cartridges, however, exhibit significant sonic differences in frequency

response, detail retrieval, spatial perspective, dynamics, and richness of tone color.

You should listen for two other aspects of turntable performance: speed accuracy and speed stability. Speed accuracy is how closely the turntable's speed of rotation approaches precisely 33 $\frac{1}{3}$ revolutions per minute (rpm). You need to worry about speed accuracy only if you have absolute pitch (the ability to identify a pitch in isolation). Speed instability, however, is easily audible to anyone, and is particularly annoying. Speed stability is how evenly the platter rotates. Poor speed stability causes wow and flutter. Wow—very slow variations in speed that shift the pitch slowly up and down—is most audible with recordings of solo piano that contain sustained notes. *Flutter* is a rapid fluctuation of speed that sounds almost like tremolo. Together, wow and flutter make the sound unstable and blurred rather than solidly anchored.

A good turntable with no obvious flutter can still suffer from speed instability. Instead of hearing flutter directly, you may hear it as a reduction in timbral accuracy—an oboe, for example, will sound less like an oboe and more like an undifferentiated tone. Speed variations also seem to erase low-level details, such as the finest information in an instrument's harmonic structure, or the tail ends of reverberation decays. In addition, tiny speed variations affect the quality described earlier: the impression of the sound being composed of individual instruments occupying specific points in space, surrounded by an acoustic.

Technical Aspects of Choosing an LP Front End

Matching a turntable, tonearm, and cartridge to each other involves not only aesthetic choices about which combinations sound most musical, but technical decisions as well.

First, the tonearm must be able to fit the turntable's arm-mounting area. Many turntables have an arm-mounting board, or armboard, to which the tonearm is fastened. The armboard must be at least as big as the arm's base, and be able to securely hold the arm—any looseness here will seriously degrade the sound. When mounted on the arm-board, the tonearm's cartridge end should be positioned within a range that allows the cartridge to be positioned at precisely the correct distance from the tonearm's pivot point, a parameter called *overhang*. Overhang (described later in this chapter, under "The Tonearm") can be set using a template supplied by the turntable manufacturer, or a third-party alignment protractor.

The turntable's suspension should be stiff enough to support the tonearm's weight. If the tonearm is too heavy for the turntable's suspension, the turntable won't be level. Turntable manufacturers specify the ranges of tonearm weights appropriate for their turntable models.

Next, the tonearm's effective mass must be matched to the cartridge's compliance. Let's define these terms before examining how they interact.

Effective tonearm mass isn't the tonearm's weight, but rather the amount of mass in combination with where along the arm that mass is located. For example, one gram of mass added at the arm's pivot point represents almost no contribution to the arm's effective mass. But one gram added at the stylus position contributes exactly one gram to the effective mass. The effective-mass specification conveys the amount of inertia the cartridge "feels" as it tries to move the arm while the stylus is being modulated by the record groove, as stylus, cartridge, and arm are all being dragged across the record's playing surface. (Technically, the arm's effective mass is the amount of mass that, if positioned at the stylus, would equal the total rotational inertia of the arm's moving parts relative to the distance from the pivot. This rotational inertia is called the moment of inertia.) A tonearm of less than 10 grams effective mass is considered a low-mass arm, 11–20 grams is considered a medium-mass arm, and more than 20 grams is a high-mass arm.

A cartridge's *compliance* describes how stiffly or loosely the suspension holds the cantilever. The *cantilever* is the thin tube that emerges from the cartridge body and holds the stylus. If the cantilever is easily moved, the cartridge is said to have high compliance. If the cantilever is stiffly mounted, the cartridge is of low compliance.

Compliance is expressed as a number indicating how far the cantilever moves when a force is applied to it. Specifically, a force of 10^{-6} dynes is applied; the cantilever's movement, measured in millionths of a centimeter, is the cartridge's compliance. For example, a low-compliance cartridge (a stiff suspension) may move only ten millionths of a centimeter; we say the cartridge has a compliance of 10. Because this method of expressing compliance is standardized, the reference to millionths of a centimeter is dropped, leaving only the value 10. Moderately compliant cartridges have compliances of 12 to 20; high-compliance cartridges are any value above 20.

A cartridge's compliance and the tonearm's effective mass form a resonant system called, in engineering jargon, a "mass-spring system." That is, the combination will vibrate much more easily at a particular frequency than at other frequencies. Think of a weight attached to the bottom of a Slinky, a popular children's toy. Imparting motion to the weight will cause it to begin cycling up and down at its resonant frequency, which is determined by the spring rate (analogous to the cartridge compliance), the amount of mass (analogous to the tonearm's effective mass), and whether the weight was bounced with a large or small motion. Similarly, a tonearm and cartridge, when put into motion, will resonate at their own resonant frequency. Energy is imparted to the tonearm and cartridge by record warps, turntable rumble, the turntable's resonance, record eccentricity (the center hole isn't exactly centered), and footfalls (the vibrations of someone's footsteps transmitted from the floor to the turntable). These energy sources are all of very low frequencies, perhaps below 8Hz, and can set the tonearm-and-cartridge system into resonance at the system's resonant frequency. The tonearm-cartridge resonant frequency must be lower than any musical signal, and higher than the resonant frequency of the turntable suspension, or great distortion of the sound can result. Although we can't avoid resonance in the tonearm and cartridge, we can adjust the latter so that their resonant frequency falls above the very low frequencies of rumble and record warp, but below the lowest musical pitch recorded in the grooves. By matching the arm's effective mass to the cartridge's compliance, we can tune the resonant frequency to fall between the sources of vibration. Fig.9-4 shows the resonant frequencies of different values of arm mass and cartridge compliance.

Another, distinctly different type of tonearm resonance is the arm itself resonating independently of the cartridge, which we'll call structural ringing. Just as a struck bell will ring at its own natural frequency, the tonearm will vibrate for a finite period of time after energy is imparted to it. Different parts of the tonearm, headshell, and armtube can ring at their own structural resonant frequencies. In good tonearm design these resonances are isolated and damped, or the arm's structure or materials are changed, to reduce ringing.




Fig.9-4 The relationship between a tonearm's effective mass, cartridge compliance, and resonant frequency. (After Merrill)

Preventing the tonearm and cartridge from ringing is of utmost importance. The audio signal is generated in the phono cartridge by the motion of the cantilever relative to the cartridge. If the arm and cartridge vibrate even slightly, that vibration is converted into an electrical signal by the cartridge. Because the

cartridge can't distinguish between groove modulations (the musical information) and tonearm resonance, distortion is mixed in with the music. Tonearm resonance distorts the music's tonal balance, colors instrumental timbre, changes the music's dynamic structure (the way notes start and stop), and destroys the sense of space and imaging on the recording.

There's an important caveat to this discussion of matching a tonearm's effective mass with a cartridge's compliance: Very often, the stated specifications are wrong. It has been brought to my attention that the specified effective mass of most tonearms is grossly incorrect. Similarly, the compliance specs of most cartridges don't describe those cartridges' true compliances. It's therefore difficult to match a cartridge and tonearm by looking at specs alone. This is where your retailer's skill and knowledge come into play; he or she will have had experience setting up hundreds of different arm-cartridge combinations over the years, and will know which pairings deliver the best sound.

The Turntable

It's easy to think of the turntable as having a minor role in a playback system's sound quality. After all, the turntable only spins the record and supports the tonearm; how much sonic influence could it have?

The answer is surprising: A high-quality turntable is absolutely essential to getting the best performance from the rest of your system. A good turntable presents a solid, vibration-free platform for the record and tonearm, allowing the cartridge to recover the maximum amount of information from the groove while minimizing interference with the audio signal.

The turntable is composed of a base, platter, platter bearing, plinth, drive system, and often a subchassis. The base is the turntable's main structure; it holds all the components, and is usually finished in black or natural wood. The platter is the heavy disc that supports the record; it rests on the bearing assembly. The plinth is the top surface of the turntable, under the platter. The drive system conveys the motor's rotation to the platter. Some turntables have a subchassis suspended within the base, and on which the platter and tonearm are mounted. Every turntable will also have an armboard for mounting a tonearm, or a slot into which the tonearm's base fits. Many turntables have no base or plinth, instead suspending the subchassis in open air.

Let's look at how these components are assembled to form the modern turntable.

The Base and Plinth

A turntable's *base* and *plinth* play important roles in sound quality. The base must be a rigid, vibration-resistant structure on which the other turntable components can be mounted. If the base is flimsy, it will vibrate and transmit that vibration to the platter and tonearm, thus degrading the sound.

A turntable system can be set vibrating by four forces: 1) acoustic energy impinging on the turntable (called feedback); 2) structure-borne vibration traveling through the turntable stand (primarily when the stand is set on a suspended floor); 3) the turntable's mechanical systems, such as the platter bearing and motor vibration; and 4) the vibration imparted to the tonearm by groove modulations.

This vibration creates relative motion between the stylus and the cartridge. Because the cartridge can't distinguish between groove modulations and turntable resonances, this vibration is converted to an electrical signal and amplified by your system. This is why turntable designers use elaborate measures to prevent vibration.

Let's first take the case of acoustic energy impinging on the turntable. If the base and plinth aren't rigid, they're more likely to be set in motion by sound striking the turntable. In extreme cases, the speakers and turntable create an acoustic-feedback loop in which sound from the speakers is converted into an electrical signal in the cartridge through vibration, which is amplified by the loudspeakers, which causes even more feedback to be produced by the cartridge, and so on. This acoustic feedback can muddy the music, or even make it impossible to play records at a moderately high playback level. You can hear this phenomenon by placing the stylus in the lead-in groove without the record spinning, then gradually turning up the volume. You'll start to hear a howling sound as the acoustic feedback loop grows strong enough to feed on itself, and a "runaway" condition develops in which the sound keeps getting louder. If you try this, keep your hand on the volume control and be ready to turn down the volume as soon as you hear the howl—if you don't turn down the volume immediately, your system could be damaged. The more resistant to vibration the turntable is, the less severe this phenomenon. You might try this with and without the turntable's dustcover in place. A turntable's dustcover can catch acoustic energy and feed it into the base and plinth, degrading the sound. This is why nearly all high-end turntables eschew dustcovers. Turntable bases and plinths are designed to resist vibration by making them so massive that it's hard

to put them in motion, and by damping any vibrations that do exist. Materials such as *medium-density fiberboard (MDF)* are often used in turntable bases. High-quality turntable bases and platters are often made from machined acrylic. Vibration damping is achieved with layers of material applied to or within the base so that the vibration stops rather than being allowed to ring. An example of vibration damping is putting a wet rag on a steel plate before striking the plate. Instead of hearing the plate ring, you hear a dull thud—the vibrational energy is dissipated more quickly into and through the wet rag than it is through the air. In practice, damping materials are often exotic compounds—one turntable maker uses the same sound-deadening material that lines a submarine’s hull. A technique called *constrained-layer damping* puts layers of soft damping material between layers of harder materials to more effectively dissipate vibrational energy. The rationale for this is that when vibration encounters a discontinuity in the material, shear forces between layers produce shear strain, which causes frictional losses of that mechanical energy as it is converted into minute amounts of heat. In another technique, lead inserts are distributed throughout the base and platter. Some designers believe that this discontinuity in materials is in fact a mistake and that, rather than absorbing vibration, it simply reflects it. These designers hold that solid blocks of an ideally chosen material, coupled with a sophisticated isolation system, deliver optimal performance.

The armboard should have properties of vibration resistance and damping similar to those of the base. Some armboards of acrylic and lead weigh up to six pounds.

Another theory of turntable design suggests that the *less* massive the base and plinth, the better. This school argues that a very rigid, low-mass base is ideal because there’s less mass to resonate.

Sprung and Unsprung Turntables

Now let's look at how turntable design can address problems caused by structure-borne vibration. Vibration entering the turntable through the stand or rack can be greatly reduced by mechanically isolating the turntable's key components (platter, armboard, and tonearm-cartridge) with a suspension system—the sprung turntable. Most turntables are *sprung*, meaning that their platter and armboard are mounted on a subchassis that floats on springs within the base. The terms *suspended* and *floating* describe the same construction.

Sprung turntables can be of one of three designs. In the first method, the subchassis sits on springs attached to the base's bottom. In the second, the subchassis hangs down from the plinth on springs. The third technique, shown in [Fig.9-5](#), dispenses with the base entirely and hangs the subchassis on pillars in open air.

Whichever technique is used, the goal of all sprung designs is to isolate the platter and tonearm from external vibration. Any vibration picked up by the supports on which the turntable rests won't be transferred as effectively to the platter and tonearm. The primary sources of structure-borne vibration are passing trucks, footfalls, air-conditioners, and motors attached to the building. Structure-borne vibration is much less of a problem in a single-family home with a concrete floor than in an apartment building, or a frame house with suspended floors.



Fig.9-5 The subchassis can also be mounted in open air, with no base. (Courtesy Basis Audio)

The suspension on which the subchassis is mounted is carefully tuned to a very low frequency. The vertical resonant frequency of the suspension can also be tuned to a frequency slightly different frequency from the horizontal resonant frequency. In addition, the suspension can be immersed in a viscous fluid to damp any remaining vibration. An advantage of viscous damping over mechanical damping such as foam or Sorbothane (a rubberlike material) is that the viscous fluid doesn't compress, thus retaining its damping properties under all conditions.

You can judge a suspension's effectiveness by putting the stylus on a stationary record and tapping on the stand on which the turntable sits. The volume should be set to what would be a moderate listening level if the record were spinning. Begin by gently tapping the stand, then gradually increase the force. The less sound you hear from the loudspeakers, the more effective the turntable's suspension. With the best models, you can even hit the stand as hard as you like with a hammer and still hear no sound from the speakers—the suspension completely isolates the platter and tonearm from the hammer blow. In less effective designs, a small tap on the stand can be heard through the speakers.

However, a minority of turntable designers believe that the springs used in suspended turntables can resonate and actually *introduce* vibration into the

platter and arm. They believe that the answer is the unsprung turntable, in which the platter and tonearm are mounted directly on the base. If the base is made rigid and nonresonant enough, they argue, the platter and arm won't vibrate—and you don't have to worry about springs resonating. Unsprung turntables sometimes use constrained-layer damping in the base in an attempt to dissipate mechanical energy before it reaches the platter and tonearm.

Whatever the design, you can judge its isolation properties by putting a stethoscope on the turntable and scratching or tapping the equipment rack on which the turntable sits. The greater the isolation, the less sound you'll hear through the stethoscope

The Platter and Bearing Assembly

The platter not only provides a support for the record, it also plays two other important roles: as a flywheel, to smooth its own rotation; and as a sink to collect and dissipate vibrations from the record. Platters can be very heavy—up to 80 pounds—with most of their mass concentrated toward the outside edge, to increase their moment of inertia and thus the flywheel effect. This high mass also counters bearing friction and stylus drag. Irregular (nonlinear) bearing friction can create rapid irregularities in the platter's speed, which causes the pitches of recorded instruments and voices to fluctuate. Massive platters are not a cure-all for poor-quality bearings—a heavy platter can't make up for the additional noise created by high bearing friction. Precise tolerances in the centering and roundness of the platter, and low friction in its bearing, are far more important than the platter's weight. Part of what you pay for in a more expensive turntable is this high-precision machining.

Most platters are made of a single substance, such as acrylic, stamped metal (in cheaper turntables), cast and/or machined aluminum (mid-level turntables), or exotic materials such as ceramic compounds. The platter sometimes has a hollowed-out ring around the outer edge that is filled with heavier materials to increase the platter's inertia, or is loaded with a damping substance to make it more resistant to vibration.

These techniques also attempt to make the platter act as a sink for record vibrations. When the record is clamped to the platter, some degree of record vibration will be transferred to the platter. How much of that energy is transferred, and over what range of frequencies, is determined by the integrity of the contact between record and platter, as well as the properties of the platter material. Vacuum hold-down systems, described later in this chapter, intimately couple the record to the platter for minimal record vibration.

Because the platter spins on a stationary object (the rest of the turntable), there must be a bearing surface between the two. With one technique, the bearing is mounted at the end of a shaft to which the platter is attached. This shaft, often made of stainless steel, occupies a hollow sleeve in the base. The shaft has some form of spherical bearing on its bottom end made from chrome-hardened steel, tungsten-carbide, zirconium, ceramic, or even a very hard jewel, such as sapphire. Often, the bearing sits in a well of lubricant. A second technique puts the spherical bearing surface atop a stationary shaft, between the platter and shaft.

Whichever technique is used, the bearing must provide smooth, quiet rotation of the platter—any noise or vibration created by the bearing will be transmitted directly to the platter. Turntable bearings are machined to very close tolerances, and are often highly polished to achieve a smooth surface.

An alternative that offers extremely low friction is the air bearing. A pump forces compressed air into a tiny gap between the platter and the plinth. The pressurized air pushes up the platter slightly so that the platter literally floats and rotates on a cushion of air rather than on a mechanical bearing. Air-bearing turntables are generally expensive and more fussy about setup.

Another method of avoiding a mechanical bearing is magnetic suspension. In such turntables, the platter and plinth contain powerful magnets (usually of neodymium) oriented to repel each other. This lifts the platter off its mechanical bearing a fraction of an inch, resulting in the platter spinning on, again, a cushion of air.

Platter Mats, Record Clamps, and Vacuum Hold-Down Systems

Platter mats are designed to minimize record vibrations and to absorb what vibrations remain. Designers of soft mats suggest that an absorbent felt mat works better in drawing vibrations away from the record. Designers of hard mats contend that a stiffer mat material better couples the record to the platter. Finally, some turntable manufacturers discourage using any mat at all, believing that their platter design provides the best sink for record vibrations. Marketing claims aside, physics dictates that the greatest energy transfer occurs when the platter's key physical properties match as closely as possible the physical properties of the record. This phenomenon of a mechanical wave being transferred from one solid to another has been well documented in other fields of engineering. Consequently, one could make a good argument that correct choice of platter material is a better solution than a record mat. Moreover, some soft mats attract and hold dust that can be transferred to the record.

Record clamps couple the record to the platter so that the record can't vibrate freely. Ideally, the platter acts as a vibration sink, draining vibrations from the record and thus improving the sound. The clamp better transfers record vibration to the platter, which then dissipates that vibration. Clamps also help to flatten warped records.

Record clamps come in three varieties. First, the clamp can simply be a heavy weight placed over the spindle and atop the record. The clamp's weight presses the record down onto the platter. Other clamps have a screw mechanism that threads down onto the spindle. The third type is the reflex clamp, in which a locking mechanism pushes the clamp down onto the record. Which type works best should be decided by your listening, the turntable manufacturer's recommendations, or your local dealer's suggestion. Note that very heavy clamps can strain some sprung turntables, compressing their suspension springs (or stretching them, if the subchassis is hung from the plinth). Another type of device designed to more closely couple the record to the platter is a heavy, 12"-diameter ring placed over the record's outer edge. This device couples the record's outer diameter to the platter, but is cumbersome to use in practice (it's difficult to center the ring on the record). Whatever route you chose, some form of record clamping is a must for any high-end turntable.

Perhaps the best way to couple the record to the platter is with a *vacuum hold-down*—the same method used in LP mastering lathes to secure the master lacquer to the lathe's turntable. In this technique, an external pump creates a vacuum that is transferred by tubes to the platter. The platter contains small channels that distribute the suction, and the platter's edge is surrounded by a flexible lip that couples to the record to form a vacuum seal. The vacuum holds the record tightly to the platter surface so that vibrations are drained away from the record. A vacuum hold-down also makes the record more flat, even if it is only mildly warped. One can easily hear the effect of a vacuum hold-down system simply by turning it off and comparing the sound. In my experience, vacuum hold-down offers significant improvements in clarity, instrumental timbres, and image focus. Detractors of vacuum hold-down suggest that the vacuum system causes dust on the platter to become embedded in the underside of the record, contributing to surface noise. I've used a turntable with vacuum hold-down for several years and have not encountered this phenomenon.

The Drive System

A turntable's drive system transfers the motor's rotation to the platter. Virtually all high-end turntables currently made are belt-driven: the platter is spun by a rubber belt or silk thread stretched around the motor pulley and outer rim of the platter.

Mass-market mid-fi turntables of the 1970s and 1980s, as well as a few current high-end models, are direct-drive; i.e., the motor is connected directly to the platter. The motor's spindle is often the spindle over which you place the record. Direct drive was sold to the public as superior to belt drive—there are no belts to stretch and wear, and a direct-drive motor can be electronically controlled to maintain precise speed and exhibit low wow and flutter. Indeed, the wow and flutter specifications of a direct-drive turntable can be better than all but the best belt-drive turntables. Direct-drives often include electronic speed control, and sometimes a *phase locked loop* (PLL) servo mechanism, both of which are claimed to ensure good speed stability. In a servo system, the platter's speed is sensed and a signal is fed back to the motor, which responds to keep the platter turning at a constant, correct speed. However, as in all servo systems, the correction signal is generated only after the system has detected a speed error, which introduces some degree of pulsation in the rotational speed. Servo systems are also prone to long-term drift, a drawback for listeners with perfect pitch. The primary advantage of direct drive is that the motor's vibrations are directly transmitted to the LP.

The vast majority of high-end turntables are belt-driven. Rather than directly couple the motor's vibration to the platter, as in direct-drive, the drive belt acts as a buffer to decouple the platter from the motor. Motor noise is isolated from the platter, resulting in quieter operation than is possible with a direct-drive turntable. Belt drive also makes it easier to suspend the platter and drive system on a subchassis.

Many belt-drive turntables do not include elaborate speed controls; the motor just sits there, spinning at a fairly high speed (as fast as 1000rpm). This high-speed rotation is coupled to the large platter with a small pulley, resulting in 331/3rpm rotation of the platter. Belt-drive motors are usually *AC-synchronous*, meaning that their rotational speed is determined by the frequency of the AC voltage driving them. But because the AC line is subject to noise and fluctuation, the platter's speed can be very slightly affected by dirty AC power. Although the resulting variations in speed are minute, it doesn't take much to destroy the

music's pitch definition, timbres, and image stability. Consequently, some manufacturers synthesize their own 50 or 60Hz AC sinewave with an outboard box. This device creates a pure, stable sinewave with a precision oscillator, then amplifies it to 120V to drive the turntable's AC synchronous motor. In one such drive system, the stock power supply produces a small but noticeable vibration of the motor. After switching to the upgraded supply, the motor runs so smoothly that one can hold it in one's hand and be unable to feel that it's running at all.

The drive motor can be a source of turntable vibrations. As the motor spins, it produces vibrations that can be transferred to the other components in the turntable, producing a low-frequency rumble. Even if you don't hear rumble directly, motor vibrations can still degrade the sound. The motor assemblies of some turntables are completely separate from the base and encased in damping material. Other designs mount the motor to the subchassis, and isolate its vibration from the other turntable components.

The Tonearm

The tonearm's job is to hold the cartridge over the record and keep the stylus in the groove. We want the tonearm to be an immovable support for the cartridge, yet also be light enough to follow the groove's inward-spiraling path, track the up-and-down motions of record warps, and follow any record eccentricity caused by an offset center hole—all without causing undue wear on the delicate grooves themselves, and all, of course, while tracing as accurately as possible the microscopic modulations of the groove walls and trough that comprise an analog of the audio signal. As we'll see, it is a challenge to any cartridge to simultaneously perform all of these jobs well.

Tonearms come in two varieties: *pivoted* and *tangential tracking*. A pivoted tonearm allows the cartridge end of the arm to traverse the record in an arc whose center is the arm's fixed pivot point. A tangential-tracking tonearm—also called a linear-, radial-, or parallel-tracking tonearm—moves the entire cartridge and tonearm across the record's surface.

Let's first look at the pivoted tonearm, by far the most popular type. Its major components are, from back to front, the counterweight, bearing, armtube, and headshell. (These elements are shown from left to right in [Fig.9-6](#).) The counterweight counteracts the weight of the armtube and cartridge; its weight and position determine the downward force of the stylus in the groove. The bearing provides pivot points for the arm in both the vertical and horizontal planes. The armtube extends the cartridge position away from the pivot point to the optimal position over the record. The headshell is attached to the end of the armtube and provides a platform for mounting the cartridge to the armtube. The small, flat disc near the bearing in [Fig.9-6](#) sets the anti-skating compensation.



Fig.9-6 A pivoted tonearm. (Courtesy Linn Products)

Bearings are required to allow the tonearm to move up and down as well as from side to side, and are important aspects of its design. The bearing should provide very low friction and not impede the arm's movement. If the bearing is sticky, the stylus will be forced against the groove wall, causing distortion and record wear. Loosening the bearing reduces friction, but can cause the bearing to "chatter" as the tonearm is rattled by the motion of the stylus in the groove, or by other sources of tonearm vibration. Remember, any movement of the tonearm in relation to the stylus in the groove is interpreted by the cartridge as a groove modulation, and is converted into an electrical signal that appears at the cartridge output as part of the musical signal. Tightening the bearings decrease chatter but also increases friction. Tonearm designers must balance these tradeoffs.

The most common types of pivoted tonearm bearing are the *gimbal* and *unipivot*. In a *gimballed* bearing, two sets of bearings are oriented at 90° to each other, much like those in a navigational gyroscope. In a *unipivot*, a hollowed-out cup in the arm rests on a point. Unipivot arms have lower friction than arms with *gimballed* bearings, but are less mechanically secure.

An interesting approach to tonearm bearings is the Well Tempered Tonearm, from Well Tempered Lab. The tonearm is attached to a vertical post and horizontal paddle suspended in a cup of silicon by monofilament line. In essence, the WTT has no bearing; the arm moves in a bath of fluid, with no hard connection between the arm and the rest of the turntable.

The Basis Audio Vector tonearm provides a twist on the unipivot arm. The Vector's innovation is to asymmetrically weight the arm (with a cutout in the counterweight) so that it "leans" over onto a second bearing. This technique eliminates the possibility of dynamic azimuth error. (Azimuth is the

perpendicular relationship between the stylus and the record groove.) Record warps can cause conventional pivoted arms to “roll,” or constantly change their azimuth alignment.

A tonearm’s armtube is designed to be rigid, low in mass, resistant to vibration, and have the ability to damp any resonances that do occur. Armtubes are often made of exotic materials that combine rigidity with low mass (carbon fiber is a popular choice), and can be filled with damping material to kill resonances. One particular arm (the Graham) is even made of concentric tubes separated by damping material. The goal is to make the arm as inert as possible, so that it is less likely to ring and transmit those vibrations to the cartridge. The better-damped an armtube is, the better it can retrieve low-level musical detail.

A tonearm’s effective mass (defined earlier in this chapter) is carefully chosen to strike a balance between the ability to hold the cartridge steady as modulations in the groove move the stylus from side to side (which requires high mass), and the ability to follow warps and eccentricities in the record (which requires low mass). When a massive tonearm encounters a record warp, the downward force of the stylus will increase on the “uphill” side of the warp, then decrease on the “downhill” side. If the warp is severe and and/or the arm is heavy enough, the stylus can even fly out of the groove. Although this problem occurs with all tonearms, it is worse with arms of high effective mass.

A tonearm’s *effective length* is the distance from the pivot point to the stylus tip. The effective length of most tonearms is somewhere between 210 and 230 millimeters (nearly 9”). (The effective length is shorter than the arm’s actual length due to the tonearm’s *offset*, the bend in the arm near the end that holds the cartridge.) The tonearm’s effective length affects the amount of *overhang*. The technical definition of overhang is the distance the stylus extends past the spindle. In general parlance, overhang means the distance between the stylus and the pivot point. Overhang can be adjusted by moving the cartridge back and forth in the headshell. Some tonearms have adjustable-length armtubes for setting overhang. These alignment procedures are described later in this chapter.

Let’s consider a fundamental limitation of pivoted tonearms that explains why the tangential tonearm was invented.

An LP is cut with a tangential-tracking mechanism: the cutting head moves in a straight line across the lacquer master. But because a pivoted tonearm traverses an arc across the record, the playback stylus rotates slightly in relation to the groove as it is swung along this arc. This rotation causes one side of the stylus to contact the groove wall slightly ahead of its nominal contact point, and the other

side of the stylus to contact the groove wall slightly behind its nominal contact point. This difference in the angular relationship of the playback stylus to the groove is called lateral tracking error, or simply *tracking error*. Tracking error increases distortion, particularly during loud passages.

More technically, tangential lacquer cutting results in groove modulation that is always at a right angle to the groove. But the back-and-forth motion of the playback stylus in response to this groove modulation deviates from a right angle to the groove's tangent line over most of the record. Specifically, the angle is typically a few degrees less than 90° at the outermost grooves. As the stylus swings in an arc toward the record center, the stylus is rotated, reducing this angular error until it reaches a point at which there's no error. At this point the playback stylus has the same geometric relationship to the groove that the cutting stylus had to the groove. As the stylus continues in its arc and rotates farther, the angular error increases in the opposite direction. The error reaches a peak and then, counterintuitively, begins decreasing until the error reaches zero at a second point on the record. The two points of zero tracking error are typically at 2.4" and 5" from the record's outer edge.

Lateral tracking error introduces harmonic distortion, primarily second harmonic, which is sonically fairly benign. Other harmonic products are present, but are lower in amplitude than the second harmonic. The sonic result is a hardening of timbre and a coarseness to the sound. The amount of tracking-error-induced distortion is a function of not only the tracking error, but also of the groove speed, which varies from the inside to outside grooves and which steadily slows as the stylus travels from the outermost to the innermost part of the groove. The amount of distortion induced by tracking error in the innermost groove is more than twice that induced in the outermost groove. That's because one revolution of the LP beneath the stylus at the outer diameter represents a linear distance about 2.5 times longer than that described by one revolution at the inner diameter. Both revolutions are completed in precisely the same amount of time. This means that the same amount of information is packed into a space less than half as long in the inner groove relative to the outer groove. A longer space means that the groove modulation can be stretched out, making it easier for the stylus to correctly track that modulation. This is why 45rpm records sound better than 33 1/3rpm records; the groove modulation is stretched out over a greater distance.

Tracking error is minimized by putting a bend in the tonearm, called the *offset*. Offset can be a sudden angle at the end of the tonearm, or a gradual bend that gives the armtube a J or S shape. The offset angle in modern 9" tonearms is

22 to 25°, which provides the least tracking error. Adding an offset to the tonearm doesn't *eliminate* tracking error; instead, it produces the least overall tracking error, with two points along the record at which the error is zero.

Various cartridge alignments have been developed to minimize lateral tracking error. The two factors that affect tracking error are the amount of offset and the overhang, or the distance between the stylus and the pivot point. The optimal alignments were worked out some time ago; E. Lofgren and H.G. Baerwald independently developed slightly different calculations of optimal tonearm/cartridge alignment geometry in 1938 and 1941, respectively. Today, these are referred to as Lofgren or Baerwald alignments; you'll see cartridge-alignment protractors that give you a choice of Lofgren or Baerwald settings.

Another way to reduce lateral tracking error is the 12" tonearm, which was once popular and has recently made a limited comeback. In theory, an arm of infinite length would have zero lateral tracking error. A 12" arm's additional length (compared with a 9" arm) reduces tracking error because the arc's degree of curvature is smaller. The 12" arm, however, pays penalties in greater mass, greater counterweight inertia (which increases record wear), and greater potential for misalignment. Small alignment errors are magnified in a 12" arm compared with a 9" arm—the longer arm's offset and overhang are reduced, increasing any alignment errors. Moreover, a 12" arm is more prone to resonances than a 9" arm, all else being equal.

The offset needed to reduce linear tracking error in pivoted tonearms introduces a problem called *skating*. Skating is a force that pulls the tonearm toward the center of the record. This makes the stylus press harder against the inner wall of the groove, and must be compensated for by applying an equal but opposite force on the cartridge. This compensation, called *anti-skating*, counteracts the skating force caused by tonearm offset. Anti-skating allows the stylus to maintain equal contact pressure with both sides of the groove, and prevents the cartridge's cantilever from being displaced from its center position in the cartridge. An anti-skating force can be generated by springs, weights with pulleys, or mechanical linkages. (I'll talk about anti-skating adjustment later in this chapter, under "LP Playback System Setup.") Skating causes uneven record wear, and a less firmly anchored soundstage.

But the problems of pivoted tonearms—lateral tracking error and skating—don't exist with tangential-tracking arms. Since the entire arm moves across the record, the angular relationship between the playback stylus and the record groove is identical to that of the LP cutting lathe at any record diameter. And because the arm has no offset, no skating force is generated. Tangential-tracking

arms have another advantage: Because they're generally very short, tangential-tracking arms can have lower vertical mass than pivoted arms. Tangential-tracking arms, however, have higher horizontal mass. Because they don't pivot, the effective horizontal mass of a tangential arm is the arm's actual mass, which is higher than the effective horizontal mass of a pivoted arm. Tangential arm design is a compromise between too-low vertical mass and too-high horizontal mass.

Tangential tracking is also much more complex, and consequently more expensive to execute well. The entire arm must be moved along a track tangential to the record—a difficult engineering challenge. In mid-fi tangential-tracking turntables, the arm is driven by a servo motor whose motion may not exactly match the record's pitch (spacing between record grooves), forcing the stylus first against one side of the groove, then the other. Although the servo system attempts to correct for this error, servo systems by their very nature work only *after* an error has occurred.

The solution in high-end tangential-tracking arms is the *air bearing*, in which the arm rides on a cushion of air surrounding a tube. The only force moving the tonearm is the gentle force of the stylus. Owing to the air bearing, the friction between the arm and the tube on which it is mounted is practically zero. (The turntable in [Fig.9-3](#) is fitted with a tangential-tracking, air-bearing tonearm.)

Tangential-tracking arms are more difficult to correctly set up, require periodic adjustment, and have air pumps that can be audible if not properly isolated from the listening area. The air rushing through the air bearing can also be audible.

Proponents of pivoted tonearms argue that the inherent geometric advantages of tangential arms are far outweighed by concerns over the tangential arm's rigidity and resonance damping. They suggest that the distortion produced by tracking error isn't a significant source of sonic degradation. Indeed, when listening to a record side all the way through, you don't feel a sense of relief at the two points of perfect tangency, nor do you notice the distortion gradually increase as the stylus leaves each of the tangency points.

The Phono Cartridge

The phono cartridge has the job of converting the modulations of the record groove into an electrical signal. Because the cartridge changes one form of energy into another (mechanical into electrical), it's called a *transducer*. There are other transducers in your system—the loudspeakers at the other end of the playback chain.

A phono cartridge consists of the cartridge body, stylus, cantilever, and generator system. The body is the housing that surrounds the cartridge, and comprises its entire surface area. The stylus is a tiny point, usually made from diamond, attached to the cantilever (the tiny shaft that extends from the bottom of the cartridge body). The stylus is moved back and forth and up and down by modulations in the record groove. This modulation is transferred by the cantilever to the generator system, the part of the cartridge where motion is converted into an electrical signal.

Moving-Magnet and Moving-Coil Cartridges

Cartridges are classified by their principle of operation: *moving-magnet* or *moving-coil*. The magnets and coils work together as an electrical generator to create an output voltage that is an analog of the relative motion between the magnets and coils. In a moving-magnet cartridge, tiny magnets attached to the cantilever move in relation to stationary coils in the cartridge body. The movement of the magnetic field through the coils induces a voltage (the audio signal) across the coils. (This phenomenon, called electromagnetic induction, is explained in detail in [Appendix B](#).) Less common variations on the moving-magnet cartridge are the induced magnet and moving-iron types. (Fig 9-7 shows the essential elements of a moving-magnet cartridge.)

A moving-coil cartridge works on exactly the same physical principles, but the magnets are stationary and the coils move. A moving-coil cartridge generally has much less moving mass than a moving-magnet cartridge. Consequently, a moving-coil cartridge can generally track better than a moving-magnet type, and also have better transient response. With less mass to put into motion (and less mass to continue moving after the motivating force has stopped), moving-coil cartridges can better follow transient signals in the record. Because of their construction, moving-coil cartridges generally don't have user-replaceable styli; you must return the cartridge to the manufacturer.




Fig.9-7 Moving-magnet cartridge construction. (Courtesy Audio-Technica U.S., Inc.)

Cartridge output voltage varies greatly between moving-magnets and moving-coils. A moving-magnet's output ranges from 2mV (two thousandths of a volt, $2/1000V$, or 0.002V) to about 8mV; a moving-coil cartridge's output is typically between 0.15 and 2.5mV. Although moving-coil cartridges generally have lower output voltage than moving-magnet types, some so-called "high-output" moving-coils have higher output voltage than some moving-magnet cartridges. A cartridge's output voltage is specified when the stylus is subjected to a recorded lateral groove velocity of 5cm per second RMS at 1kHz.

This wide range of cartridge output voltage requires that the phono preamplifier's gain be matched to the cartridge's output voltage. The lower the cartridge output voltage, the higher the gain needed to bring the phono signal to line level. (A full discussion of matching cartridge output voltage to phono stage gain is included in [Chapter 4](#).)

A moving-coil's output voltage is determined largely by the number of turns of wire on the coil: The more turns, the higher the output. More turns means less gain is needed in the phono preamp, but most designers use the minimum number of turns possible to keep the moving mass low.

The method of generating the audio signal, whether moving-magnet or moving-coil, affects how sensitive the cartridge is to the effects of preamplifier "loading." The cartridge output "sees" an electrical load of the preamplifier's input impedance and capacitance plus the tonearm cable's capacitance. The load impedance is in parallel with the capacitance.

The frequency response of a moving-magnet cartridge is greatly influenced by the amount of capacitance in the tonearm cable and preamplifier input. A moving-magnet cartridge terminated with the correct impedance (47k ohms) but with too much capacitance, can have frequency-response errors in the upper midrange and treble of as much as 5dB. Most cartridge manufacturers specify the ideal load impedance and capacitance for their cartridges. A typical load for a moving-magnet cartridge is 47k ohms in parallel with 200–400 picofarads (pF) of capacitance. Moving-coil cartridges are nearly immune to capacitance-loading effects.

The Strain Gauge Cartridge

Not all phono cartridges work on the principle of electromagnetic induction. The strain gauge cartridge, first introduced in the 1960s and making a very limited comeback today, translates an LP's groove modulations into an electrical signal using a completely different principle. In the strain gauge cartridge, the cantilever is mounted between two silicon semiconductors whose electrical resistance changes in response to mechanical pressure. A special preamplifier designed specifically for the cartridge supplies continuous current flow through the semiconductors. The LP's groove modulations cause the cantilever to press against the semiconductor crystals, varying the current flow through them, and this variation in current flow represents the audio signal. Note that all strain gauge cartridges are sold with a preamplifier; the two components are integral to the system's operation.

Stylus Shapes and Cantilever Materials

Styli come in a variety of shapes, the simplest and least expensive of which is the conical or spherical tip. The conical stylus is a tiny piece of diamond polished into a rounded point. Because the stylus tip must be smaller than the width of the record groove, the radius of a conical stylus tip is about $15\mu\text{m}$, or about $0.0006''$.

A conical stylus, however, can't perfectly track the groove's modulations because its shape is different from that of the cutting stylus used to make the record, which has a more elliptical cross section. As shown in Fig.9-8, a conical stylus contacts points on the groove wall different from those contacted by the cutting stylus, causing tracing distortion. In a related problem, called *pinch effect*, the stylus is pushed up out of the groove by the groove walls. This occurs because the groove width isn't constant at all points of the groove modulation. If you imagine a sinewave recorded on an LP, the groove is narrower at the sinewave's zero-crossing points, a result of the LP having been cut with a triangular-shaped stylus. When the conical stylus encounters the narrower portions of the groove, it is pushed up slightly.

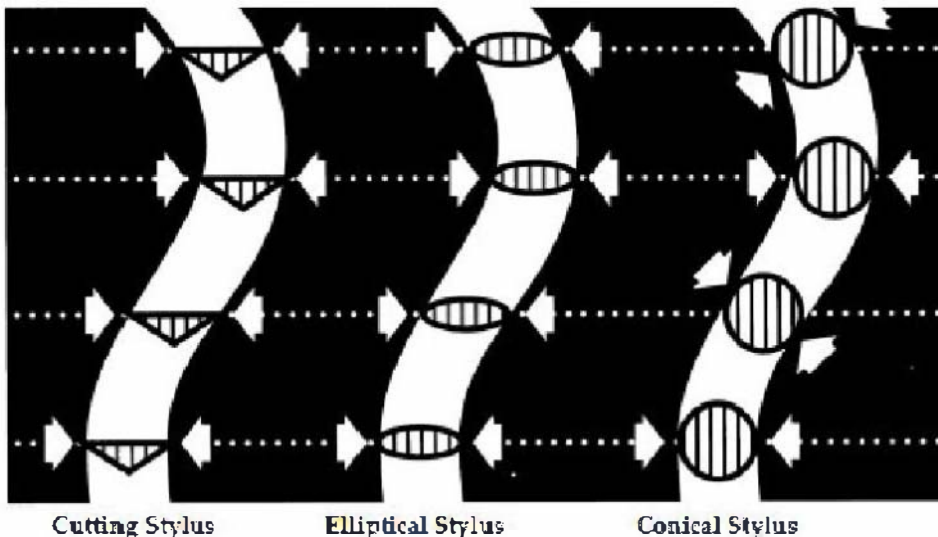


Fig.9-8 Comparison of conical and elliptical styli in relation to the cutting stylus (After Shure)

The solution to these two related problems is the elliptical stylus. Instead of using a rounded tip to track the groove modulations, an elliptical stylus has an oval cross section, with two flattish faces. Because this shape more closely

approximates the shape of the cutting stylus, it results in lower tracing distortion and eliminates the pinch effect. An elliptical stylus can also better track high-frequency modulations, and distributes its force over a wider area to reduce record wear. (Fig.9-8 compares the stylus/groove relationships of cutting, elliptical, and conical styli.)

A third type of stylus takes the elliptical shape one step further by making the tip even narrower. These stylus shapes are called fine-line, hyperelliptical, line-contact, Shibata, and van den Hul. A cartridge with any of these stylus shapes requires more precise alignment and setup than do conical or standard elliptical styli.

A variation of the elliptical stylus is the microridge, which has a tiny groove cut into each edge. As a microridge stylus wears, it retains its shape and thus its relationship to the groove.

As described later in this chapter, keeping your stylus clean is of paramount importance. A speck of dust or dirt attached to the stylus is like a boulder grinding away at the groove walls. An appreciation of the enormous pressure a stylus imposes on the groove further highlights the need for a clean stylus. For example, a tracking force of 1.4 grams applied to a typical stylus contact area (0.2×10^{-6} ”) results in a pressure of nearly four tons per square inch—enough to momentarily melt the outer layer of the groove wall. It’s easy to see how stylus motion through the groove is much smoother with a clean stylus, and produces much less record wear. A clean stylus sounds better, too.

With good maintenance, a stylus should last for about 1000 hours of use. It’s a good idea to have the stylus examined microscopically after about 500 hours, then again at 800 hours to check for irregular wear that could damage records.

Because the cantilever transfers the stylus’s movements to the generator, its construction is extremely important. Cantilevers are designed to be very light, rigid, and non-resonant. The lower the cantilever mass, the better the cartridge’s ability to track or trace the groove, all else being equal. Mistracking occurs when the stylus can’t accelerate quickly enough to follow the groove modulations, and the stylus momentarily loses contact with the groove wall, resulting in tracking distortion. A stiff, low-mass cantilever contributes to good tracking ability. To obtain stiffness with low mass, cantilevers are often made of exotic materials, including boron, diamond, beryllium, titanium, ceramic, ruby, and sapphire. Cantilevers are often hollow to reduce their mass, and are sometimes filled with a resonance-damping material. Most test records include bands of increasing groove velocity that allow you to assess a cartridge’s tracking ability. A famous

example of an LP that caused virtually all cartridges to mistrack is the Telarc recording of the *1812 Overture*; the groove modulation representing the cannon firing would cause some cartridges to not just mistrack, but to jump off the record's surface.

The cantilever is mounted in a ring or block made of a compliant rubber, called the damper, inside the cartridge body at the end opposite that which bears the stylus. The damper allows the cantilever to move while keeping it in position.

Some audiophiles remove the cartridge body to run the cartridge "naked." The body can be a source of resonance, and adds a fair amount of effective mass to the tonearm. Removing a cartridge body should be attempted only if you are very skilled in working with tiny precision devices, and then only if you're willing to risk voiding the cartridge's warranty—and destroying your cartridge.

LP Playback System Setup

Correctly playing back a record is a delicate art. It takes patience, skill, a keen ear, a delicate touch, an appreciation for the forces involved—and a high-quality LP front end. Doug Sax, co-founder of the record label Sheffield Lab and the pioneer of modern direct-to-disc recording, has said, “Probably the easiest thing in the world is to play an LP record incorrectly.”

Sax’s observation was prompted by the great variability in setup possible with a turntable, tonearm, and cartridge. When you add to the equation how the phono preamp loads the cartridge, the turntable stand, and its placement in the room, you have the potential for not getting all the music out of the record grooves.

But this situation also offers the promise of improving the sound of your LP front end without spending any more money. Just by setting up the system correctly, you can realize far better sound from a modest turntable than from a poorly set-up front end costing thousands more. Great musical rewards await those willing to tweak their LP front ends.

The best sources of advice for setting up an LP playback system are the manufacturers of the components in the LP system, and your dealer. The dealer will have undoubtedly set up many turntable-tonearm-cartridge combinations like yours, and will have the techniques down pat. A good dealer will set up your system when you buy a product from him.

Nonetheless, you should know the procedure for getting the best sound from your LP front end. This will let you converse with the dealer from a position of knowledge, and also allow you to set up your system if you just change cartridges or move your turntable from one location to another. You might also consider the excellent DVD *21st Century Vinyl: Michael Fremer’s Practical Guide to Turntable Set-Up*. Fremer walks you through all the steps in optimizing an analog front-end. Another excellent visual demonstration of cartridge and tonearm setup is the DVD that’s supplied with the Basis Vector tonearm.

The first rule of LP system setup is to put the turntable on a good platform. I can’t emphasize this enough; a flimsy rack, or one prone to resonating, will seriously degrade the sound of your system. The less good the turntable’s suspension, the more important the rack. The very best turntables—those having superlative isolation from structure-borne vibration—are very little affected by the stand. With all but the very best turntables, you won’t hear your system at its

best without a solid, nonresonant support structure for your turntable. Many stands and racks are made specifically for turntables. Invest in one.

The most massive racks can, however, put too much stress on a suspended wood floor. In addition, footfalls and other vibrations can be transmitted right through the rack and to the turntable if the turntable has inadequate suspension. If you live in a building with a shaky floor, you have several options. The first is to use a turntable shelf that mounts to the wall instead of to the floor. The second option is to place the rack directly above the point in your room where the floor joists meet the foundation. The third alternative, which can be used in conjunction with the second, is to support the joists with house jacks. Put a concrete block on the ground below the floor joists, set the house jack on top of the block, and raise the jack to meet the joist. Continue raising the jack about half a turn past where contact is made between jack and joist. Don't try to raise the joist; all you want to do is prevent it from moving up or down. You can further increase the turntable's isolation from vibration by using a commercially available product that decouples the turntable from the stand. Finally, the best solution is to buy a turntable with a superior suspension system.

Once you've got a good rack, where you place it in the room can have a large effect on the sound. Here's why: When the air in a room is pressurized by sound from the loudspeakers, the air pressure isn't evenly distributed throughout the room. Instead, there are stationary pockets of high and low pressure, called standing waves. You can demonstrate this by playing the same brief passage from a pipe-organ recording in a constant loop and walking around the room; you'll hear the sound get louder and softer. Moreover, the pressure tends to build up along walls and in corners. If the turntable is in an area of high pressure, it will be more prone to resonate from the acoustic energy impinging on it. Conversely, a turntable set up in a room "null" (a stationary area of low pressure) won't be as affected by sound from the speakers striking it. You can avoid the worst turntable placements by following a few simple rules. First, don't put the turntable in a corner, where bass tends to build up. Second, keep the turntable away from the wall, if possible. An extra 8" or 12" can make a difference.

With the turntable on your stand, make sure that the turntable is level by using a pair of small round bubble levels and adjusting the turntable's feet. If you have a sprung turntable, make sure the tonearm cable doesn't impede the motion of the subchassis.

I'll assume you've followed the manufacturer's directions for installing the tonearm and are ready to mount the cartridge. After you've attached the four tiny tonearm-cable clips to the cartridge pins, mount the cartridge in the headshell

with minimal torque on the bolts—you'll be adjusting the cartridge's position in the headshell later. Just before attaching the tonearm cable clips to the cartridge, consider applying a contact enhancer to the cartridge pins, such as DeoxIT, from Caig Laboratories.

Move the counterweight forward or backward until the arm floats, evenly balanced; then move the counterweight slightly forward, to apply some tracking force. Use a stylus-pressure gauge, also called a *vertical tracking force (VTF)* gauge—to set the VTF (Fig.9-9) to the cartridge manufacturer's specification. (Some tonearms include built-in calibrated VTF gauges.)



Fig.9-9 A vertical tracking force gauge (Courtesy AcousTech)

The next step roughly sets the *vertical tracking angle (VTA)*, also called the stylus rake angle (SRA). Adjusting the VTA changes the angle at which the stylus enters the groove. Ideally, the VTA of the playback stylus should match the VTA of the cutting stylus. Usually, this angle is correct when the bottom of the cartridge's body is exactly parallel with the record surface (Fig.9-10a). VTA is adjusted by moving the rear of the tonearm up or down at the pivot point. Negative VTA is when the rear of the cartridge is below the level of the front of the cartridge (Fig.9-10b). Positive VTA is when the cartridge's rear is above the level of the front (Fig.9-10c). Follow the cartridge manufacturer's suggestions for rough setting of VTA. (We'll fine-tune the VTA later.)

A precision protractor is essential to correctly setting the overhang and the offset, in order to minimize lateral tracking error (as we saw earlier). Overhang is adjusted by moving the cartridge forward and backward in the headshell, and offset by twisting the cartridge body in the headshell. The templates sometimes provided with tonearms are accurate only if the templates are designed for the specific tonearm, and include an outline of the tonearm printed on the template.




Fig.9-10 Vertical tracking angle (VTA) adjustments

Fig.9-11 shows a high-quality protractor that offers both Löfgren and Baerwald alignments. A quality protractor is considerably more expensive than a paper template, but much more accurate. Note that all alignment devices assume that the cartridge's cantilever is parallel with the cartridge body, which isn't always the case (see next paragraph).



Fig.9-11 A high-quality protractor allows you to optimally set up your tonearm and cartridge. (Courtesy Dr. Feickert Analogue and AcousTech)

Azimuth describes the cantilever's perpendicular relationship to the record groove. A cantilever with perfectly aligned azimuth will form a right angle with the record (Fig.9-12b). Notice that I'm talking about the *cantilever's*—not the

cartridge's—relationship to the groove. Many cantilevers are slightly skewed in their mountings, making it impossible to achieve perfect azimuth alignment by looking at the cartridge body alone. Instead, you can put the stylus tip on a small mirror and look for a perfectly straight vertical alignment of the cantilever and its reflection. If there's a bend in the reflection at the mirror, the azimuth is wrong. Some tonearms are designed for ease of azimuth adjustment; those that aren't require you to insert tiny shims between the headshell and the cartridge. Azimuth alignment is vital for correct left/right channel balance, and thus accurate soundstaging and imaging. Azimuth can also be checked with a dual-trace oscilloscope and a test record. With each channel feeding one input of the 'scope, play a test record with a pure tone. Perfect azimuth alignment will be displayed as identical 'scope traces. A simpler method for ensuring perfect azimuth alignment without an oscilloscope is the Fosgate Fozgometer, a \$250 device that, when used with a test record, allows you to quickly dial in perfect azimuth alignment.




Fig.9-12 Azimuth is the perpendicular relationship of the cantilever to the record.

Recheck the cartridge alignment and overhang with the protractor; when these are correct, gently tighten the cartridge mounting bolts. The tightness of the bolts should strike a balance between making the cartridge tight in the headshell and distorting the cartridge body. Don't overtighten them.

Following the tonearm manufacturer's instructions, set the anti-skating force. If the tonearm doesn't have a calibrated anti-skating scale, you can set the anti-skating with a grooveless record (test pressings generally have one grooveless side), or a sheet of Mylar cut into the shape of a record with a spindle hole. Put the stylus about halfway between the spindle and the outer edge, and spin the

grooveless record or Mylar sheet. Adjust the anti-skating so that the arm moves very slowly toward the outside of the record. If the anti-skating is set incorrectly, mistracking of loud passages will be more audible in one channel than in the other, and this distortion will be audible in one channel before it is audible in the other channel. Specifically, too much anti-skating will cause the left channel to distort first; too little will make the right channel distort first.

Another trick for setting the anti-skating force is to sum the left and right channels to mono and invert the polarity of one channel. When playing a record, anti-skating will be properly set when the least sound is heard from the loudspeakers. Correct anti-skating adjustment will produce equal output from both channels; when one channel's polarity is inverted, the greatest cancellation (i.e., the least sound) will occur when the output levels are most closely matched. To make an inverting cable, buy a Y-adaptor with two female RCA jacks and one male plug. Cut off one female jack, and reverse the wire's shield and center conductor to invert that lead's polarity. The turntable output drives the two female jacks, and the single output feeds one phono input on the preamplifier. This method works very well, and costs about \$3 for the Y-adaptor.

The next step is to fine-tune the VTA by listening. Mark on the tonearm rear the position at which the cartridge bottom and the record surface are parallel. Raise the back of the tonearm a few millimeters, then listen. Lower the back of the tonearm a few millimeters below the position at which the cartridge bottom and LP surface were parallel and listen again. Find the position that provides the smoothest tonal balance and the best imaging. Use several different records for setting VTA, not just one or two.

There are many more ways to improve the sound of your LP front end. You can: replace the stock tonearm cable with a high-end one; place isolation feet under the turntable; put contact enhancer on the cartridge pins; or put damping material on the tonearm tube, cartridge, finger lift, and counterweight(s). Damping material should be applied with care; it may upset the careful balance and resonance tuning designed into the turntable and tonearm.

You can check the platter's speed with a strobe disc and light. Hold the light above the strobe disc as it spins on your turntable. When the platter is spinning at precisely the right speed, the dots on the disc will appear to be standing still. Some outboard motor drives include adjustments for fine-tuning platter speed.

Another way of getting the best performance from your analog front end is to hire a local turntable-setup expert. Such experts often work in high-end retail stores; others are independent but can be recommended by your dealer.

Watching such an expert fine-tune your 'table can also be a great learning experience.

Record Care and Cleaning

How well you care for your records and stylus will greatly affect the sound quality of your analog front-end. Many of the sonic defects the general public associates with LPs (ticks, pops, skips) are the result of mishandled records. I have records that I bought in the early 1970s and have played countless times that are still in good condition. Picking up a few simple habits will allow your record collection to sound its best, as well as deliver decades of musical pleasure. For a thorough visual presentation on record care, I recommend *It's a Vinyl World, After All: Michael Fremer's Guide to Record Cleaning, Storage, Handling, Collecting & Manufacturing in the 21st Century*. This DVD also takes you behind the scenes at an LP mastering session, and shows you how LPs are made.

Records should be stored vertically and packed tightly enough so that they remain at a right angle to the floor (not leaning), but not so tightly that they are pressed together. You've no doubt seen LP jackets with a white ring in the center—the pressure from too-tight storage creates a raised area around the label that wears away more quickly than the rest of the jacket. The pressure can force particles from the paper sleeve into the grooves, where they are heard as pops and ticks. Records should also be stored away from heaters and direct sunlight. At a moderately high temperature, a record left leaning at an angle only slightly less than 90° will lose its flatness. You can buy an LP flattener that gently heats a record and applies an even pressure across its entire surface. It looks like a waffle iron, but with flat interior surfaces.

The most important rule of record care is to never touch the playing surface, even on the outside edge. Instead, remove the record and sleeve together from the jacket, then slide your hand into the sleeve until your fingers can touch the label without touching the playing surface. Remove the record with your middle finger secured by the spindle hole, and a finger touching the label on either side of the middle finger, and your thumb on the record's edge. When the record is clear of the sleeve, hold the record by its edges with both hands. I've seen people who should know better (manufacturers at trade shows, for example) pull a record out of its sleeve by grasping the record's lead-in groove area. This smears body oils on the outermost grooves, which the stylus then distributes throughout the rest of the record. Even worse, the deposited oil then acts to hold dirt and debris on the record and stylus.

It's a good idea to replace paper and plastic record sleeves with sleeves made of rice paper. Paper sleeves leave tiny fibers on the disc surface, and plastic sleeves leach chemicals as well as disintegrate over time, leaving bits of themselves on your records. Because rice paper isn't a wood-based product it doesn't shed fibers.

Before playing an LP, remove dust from the turntable with a microfiber cloth. After laying the LP on the turntable, a quick clean of the record with a carbon-fiber brush removes surface dust. Hold the brush gently above the spinning record for a few revolutions.

Ideally, the stylus should be cleaned before playing each record side. A stylus picks up dirt as it is dragged through the record groove, which alters its shape and thus degrades the sound quality and accelerates groove wear. Cleaning the stylus with a commercially available brush and a liquid cleaning solution will remove this dirt. A stylus should be cleaned with a back-to-front motion with strokes that begin at the rear of the cantilever and end by brushing the stylus itself. Some cartridge manufacturers recommend that no cleaning fluid be used; others suggest that a fluid is essential to removing accumulated dirt. There is also debate over the best type of stylus brush. Some have short, stiff bristles; others resemble nail-polish brushes. Your best bet is to follow the cartridge maker's cleaning instructions.

Dirt isn't the only stylus contaminant; the friction between the stylus and groove heats the stylus enough to slightly melt the vinyl of the groove walls. Some of this melted vinyl then sticks to the stylus, with obvious detrimental effects. This melted-on vinyl can't be removed with most commercial stylus cleaners, but there's a simple and free technique that removes even the most stubborn gunk from your stylus. Hold a matchbook from which all the matches have been removed under your stylus and scrape the abrasive striking area under the stylus in a back-to-front motion (in the same direction as the record moves under the stylus). This abrasive action removes the accumulated coating of vinyl on the stylus. This procedure must be performed with great care to avoid ripping the cantilever from the cartridge.

The coils of moving-coil cartridges become slightly magnetized over time by their proximity to the cartridge's fixed magnets, and this degrades fidelity. The solution is to demagnetize the cartridge with a small battery-operated device that connects to the tonearm's RCA cables. Some phono stages have a built-in demagnetization circuit: you simply press a front-panel button and your cartridge is demagnetized. Note that moving-magnet cartridges don't need demagnetizing, and can be ruined by trying to demagnetize them.

This routine for playing each LP should be augmented with occasional deep cleaning of records with a heavier-duty brush and cleaning solution, or with a record-cleaning machine. If you don't want to invest in a full-bore vacuum record cleaner, you can use a cleaning solution and a brush, such as the one shown in Fig.9-13. Apply the fluid, scrub the record, then wipe off the record and let it dry (in, say, a dish rack). A better approach is to use a cleaning machine with vacuum action. After placing the LP on the machine's turntable, apply fluid and allow the machine's integral brush to scrub dirt out of the grooves as the record spins. (It's a good idea to spin the record in both directions during this cycle.) You then move a vacuum arm over the record and allow the vacuum system to suck up the fluid, in which any dirt is suspended. (You should periodically clean a machine's brushes with another brush and record-cleaning fluid.) You must also keep the machine's turntable clean. If you clean one side of a record, then turn that record over to clean the other side, the cleaned side will be in contact with the turntable, potentially picking up any dirt the first side has left on the turntable. Avoid felt mats on cleaning machines; they are dust magnets and hard to clean.

Cleaning a record with a good vacuum system can render amazing results. The surfaces get much quieter, with a blacker background. Even new records benefit from cleaning. Not only do new records have dust on their surfaces that was picked up at the pressing plant, they are coated with mold release—a chemical in the raw vinyl that rises to the surface as the record is pressed, allowing the record to be more easily removed from the press.

The ultimate in LP cleaning is realized with an *ultrasonic cleaner* (Fig.9-14). These devices use ultrasonic sound waves rather than mechanical action to dislodge debris from the grooves. Ultrasonic cleaning is routinely used in the electronics, jewelry, and medical industries, and is more effective than cleaning LPs with a brush because the features within an LP groove are much smaller than even the finest bristles. Consequently, the bristles can't reach into the groove's tiny crevices. The sound waves created by an ultrasonic cleaner have no such restriction, loosening dirt anywhere in the groove.



Fig.9-13 A brush, cleaning solution, and replacement sleeves are essential parts of a high-quality LP-playback system. (Courtesy Music Direct)

In practice, ultrasonic cleaners are built around a tank of water into which a portion of the LP is immersed. The LP slowly rotates within the tank as the ultrasonic cleaning action goes to work. Unlike most LP cleaners, which use a detergent or chemical cleaning agent, the solution in the ultrasonic cleaner is simply distilled water. Some ultrasonic cleaners are supplied with a small bottle of liquid that you add to the distilled water that reduces the water's surface tension to enhance the cleaning effect.

After the LP has rotated within the cleaning bath for about three minutes, the tank drains and a blower is activated to dry the LP. The entire process is automated, and takes about five minutes per LP. With less expensive ultrasonic cleaners, you remove the LP from the bath and allow the LP to air dry. Some ultrasonic cleaners feature larger tanks that can clean as many as six LPs at a time.



Fig.9-14 An ultrasonic record-cleaning machine. (Courtesy Audio Desk)

The sonic benefits of ultrasonic cleaning are as remarkable as they are surprising. In addition to quieter surfaces, ultrasonic cleaning increases the resolution of low-level detail, deepens soundstages, and produces a greater impression that the sound is detached from the loudspeakers.

Records can also become electrostatically charged, which causes them to attract and hold dust. You can reduce this phenomenon with a device that removes the static charge. One such antistatic device looks like a gun; squeezing the trigger puts pressure on a piezoelectric crystal, causing it to generate an electric field.

It might sound like a lot of work to properly care for records and your stylus, but once these processes become habits, they will seem second nature, and you'll be rewarded with better-sounding LPs that can last a lifetime.

Vinyl as Art: Half-Speed Mastering, 45rpm Pressings, 180-gram Vinyl, and Direct-to-Disc LPs

The resurgence of vinyl playback has fueled a demand for the highest-quality records, and many specialty companies have risen to the challenge, creating LPs of exceptional sound. These endeavors are often labors of love, with the proprietors painstakingly researching the provenance of the recordings to find the true original master tapes, digging up the original photographs of classic recording sessions, creating the LP with no-holds-barred technology, and presenting an album in better-than-original sound and packaging. We'll look at several of the special techniques used in making today's audiophile LPs.

The first technique is half-speed mastering, in which the master tape is played and the mastering lathe's turntable is spun at half their normal speeds. If the original tape was recorded at 30 inches per second (ips), the tape is played at 15ips to cut the record. Concomitantly, the turntable spinning the master lacquer runs at 162/3rpm rather than 331/3rpm. The net result is that the record, when played at 331/3rpm, will reproduce the proper pitches. One advantage to half-speed mastering is that because the signal is halved in frequency (e.g., 20kHz becomes 10kHz) and the cut groove's linear velocity is also halved, the cutting stylus has more time to precisely cut the signal into the master lacquer. In addition, the amplifier driving the cutting head need produce only half the power, because amplifier power is proportionate to stylus velocity. The cutting amplifier isn't stressed, and less heat builds up in the cutting head's coils. Of course, it takes twice as long to cut a half-speed lacquer, but the effort is worth it; half-speed-mastered LPs sound better than conventional LPs, all else being equal. (Fig.9-15 shows an LP mastering lathe.)

Some records are reissued on 12" LPs cut at 45rpm. Playback requires a turntable that will spin at 45rpm (virtually all do). Because records cut at 45rpm have shorter playing times, a regular 40-minute record is often released on two discs. The advantage of cutting and playing back a disc at 45rpm is significant. With the master lacquer moving faster under the cutting stylus, wavelengths become longer and are easier to cut and to play back. The cutting stylus has more time to cut the signal, and the playback stylus more time to trace the groove modulation. Of course, 45rpm records are more expensive because two discs are required for an album that fit on one 331/3rpm disc.



Fig.9-15 A pair of LP mastering lathes. (Courtesy The Mastering Lab; photo by Neil Gader)

You've probably noticed that the inner grooves of an LP don't sound as clean, open, and extended as the outer grooves. The reason is that the recorded wavelengths become shorter toward the center of the record. Because the record spins at the same speed regardless of where the stylus is (a format called *constant angular velocity*), each revolution of the record when the stylus is tracing the outermost groove travels a much greater distance than when the stylus is tracing the innermost groove (because the distance around the outermost groove is 2.5 times the distance around the innermost groove). Toward the center of the record, the signal gets "scrunched up." High frequencies, in particular, are affected by this phenomenon because they have the shortest wavelengths to begin with. A 45rpm record ameliorates this phenomenon by making all wavelengths longer.

Incidentally, the CD is a constant-linear-velocity format: the disc's speed of rotation varies according to the playback radius. The CD spins at about 500rpm

near the innermost radius and at 200rpm at the outermost radius so that the laser can read all of the digital data at the same speed (constant linear velocity).

Most audiophile LPs today are pressed on 180-gram or even 200-gram discs. This makes for a thicker, more solid record that is less prone to vibration. Today's premium LPs are also manufactured to a much higher standard than LPs made when records were mass-market items. For example, the record is held in the press longer, which ensures that the molten vinyl completely fills the void between the stampers. You might have noticed that if an LP is noisy, the noise usually occurs at the outermost grooves. This is because the small vinyl "biscuit" that is pressed into a record flows outward from the spindle hole, toward the edges of the stamper; a too-short press cycle doesn't allow the vinyl to fully fill in the outer edges of the record, resulting in surface noise. With LPs now selling at premium prices, record manufacturers can lavish on these records longer press-cycle times, as well as 180-gram premium-quality vinyl. Premium records are sometimes made from "virgin vinyl"—that is, vinyl that has not been recycled, as was a common practice.

The highest-quality LP results from direct-to-disc recording, in which the master lacquer is cut in real time with the signal from the mixing board as the musicians perform, with no intervening tape or digital storage. Let's consider a little history before getting to modern direct-to-disc LPs.

The first sound recording, made 137 years ago, was created by cutting into a piece of tinfoil a physical representation of the acoustic waveform. This direct-to-disc process (actually, in the case of Edison's Talking Machine, direct-to-cylinder) was the only available recording method for many decades, until the invention of the wire recorder and, later, the tape recorder. The tape recorder freed musicians from the burden of performing the entire side of a record live in the studio. Each track could be recorded individually, edited into an album, and later cut into a record. The tape machine also divided the recording, editing, and disc-mastering processes into separate events. Before the widespread use of tape, which began in 1948, a musical group would go into the studio in the morning and have a finished record later that day.

With all of the advantages of tape, recording to a disc was happily abandoned by musicians and engineers, despite its sonic superiority. The world forgot about direct-to-disc recording until 1968, when Doug Sax and Lincoln Mayorga created the record label Sheffield Lab. Their stupendously great-sounding LPs demonstrated, beyond any doubt, that the purest recording format extant was the direct-to-disc LP. Sax is known today as the father of modern direct-to-disc

recording and of stereo direct-to-disc recording (before the advent of tape, all disc recording was mono).

A few other purist labels followed Sheffield's lead and made direct-to-disc recordings, but the process was largely abandoned because of the immense difficulties involved. Making a direct-to-disc recording isn't for the faint of heart. In conventional LP mastering, the physical cutting of the record—the last creative step and the first manufacturing step—takes place in an environment of perfect calm. The mastering engineer and the record's producer work with a finished master tape, transferring the electrical signal from the tape machine to the physical motion of the lathe's cutting head. If something goes wrong, or the results aren't optimal, they simply put on a new lacquer and start over. There are no musicians, no microphones, no mixing, and no performances to worry about.

Contrast that scene with a direct-to-disc session. The band plays together in the studio, with no opportunity of overdubbing less-than-perfect parts or "fixing it in the mix." The group must not only play together, but perform the entire LP side with no mistakes. Once the cutting head begins cutting the lacquer at the outside of the record, there's no turning back.

The challenge for the recording engineer is to mix all the microphones to stereo on the fly, as the band plays. If he misses a cue—bringing up the microphone on a soloist, for example—the band has to start over and play the entire record side again. For the mastering engineer, cutting the LP in real time as the band plays, the challenge is to cut that signal into a groove with the highest possible sound quality but with no technical errors. In cutting a disc from tape, an additional playback head (called the preview head) picks up the signal before the playback head. This preview signal is used solely to control the groove spacing (called the pitch). When the preview head picks up a high-level signal, it spaces the grooves farther apart, to accommodate the greater excursions of the cutting head that will occur a fraction of a second later. When the signal is low in level, the preview system packs the grooves more tightly together. The invention of variable pitch, along with the slower rotational speed to 33 $\frac{1}{3}$ rpm, is what made the long-playing record possible. The direct-to-disc mastering engineer enjoys no such safety net. Because the signal is coming right off the recording console, there's no preview, which means the mastering engineer must constantly adjust the pitch by hand during the entire LP side, based on his anticipation of the music's dynamics and bass content. If the average pitch is too wide, he'll run out of room on the lacquer. If he fails to anticipate a high-level signal, the grooves will crash into each other ("overcutting") and ruin the take.

Getting every element of the musical performance, engineering, and disc-cutting absolutely right—all at the same time—is a daunting challenge. But consider the sonic advantages of removing from the signal path a tape machine and all its electronics, not to mention the second pass of the signal through a recording console. Direct-to-disc recording not only greatly shortens the signal path for better sound, it has enormous implications for the entire recording process and, consequently, the musical result. Musicians in a typical multitrack recording session listen to the rhythm tracks through headphones and play their parts—sometimes dozens of times—until they're right. Often, the musicians recording a given track weren't even in the studio at the same time. It can be a soul-sapping experience.

A direct-to-disc session couldn't be more different. The band not only plays together, but shares the adrenaline rush of knowing that this is the take that's going on the record. It's a transformative experience. The result is a musical energy that just doesn't exist in a multitrack session.

For the music lover and audiophile, there's no closer experience to a live musical event than a direct-to-disc LP played on a first-rate analog front end.

10 Cables and Interconnects

Introduction

Loudspeaker cables and line-level interconnects are an important but often overlooked link in the music-playback chain. The right choice of loudspeaker cables and interconnects can bring out the best performance from your system. Conversely, poor cables and interconnects—or those not suited to your system—will never let your system achieve its full musical potential. Knowing how to buy cables will provide the best possible performance at the least cost.

In this chapter we'll look at all aspects of loudspeaker cables and interconnects. We'll cover balanced and unbalanced lines, bi-wiring, matching cables to your system, and how to get the most cable for your money. Moreover, we'll see why the most expensive cables and interconnects aren't always the best.

But first, let's start with an overview of cable and interconnect terms.

Cable: Often used to describe any wire in an audio system, *cable* more properly refers to the conductors between a power amplifier and a loudspeaker. Loudspeaker cables carry a high-current signal from the power amplifier to the loudspeaker.

Interconnect: Interconnects are the conductors that convey line-level signals in an audio system. The links between source components (turntable, CD player, tuner, tape deck) and the preamplifier, and between the preamplifier and power amplifier, are made by interconnects.

Unbalanced Interconnect: An unbalanced interconnect has two conductors and is usually terminated with RCA plugs. Also called a *single-ended* interconnect.

Balanced Interconnect: A balanced interconnect has three conductors instead of two, and is terminated with three-pin *XLR* connectors. Balanced interconnects are used only between components having balanced inputs and outputs.

Digital Interconnect: A single interconnect that carries a stereo digital audio signal in the S/PDIF format, usually from a CD transport, music server, or other digital source to a DAC.

USB Cable: A single cable that carries digital audio from a computer-based music server to a USB-capable DAC. (Digital interconnects and USB cables are described in depth in [Chapters 7 and 8](#).)

FireWire Cable: Also called IEEE1394, FireWire is a bidirectional interface that can carry high-resolution digital audio.

Bi-wiring: Bi-wiring is a method of connecting a power amplifier to a loudspeaker with two runs of cable instead of one.

RCA Plug and Jack: RCA plugs and jacks are the most common connection termination for unbalanced signals. Virtually all audio equipment has RCA jacks to accept the RCA plugs on unbalanced interconnects. RCA jacks are mounted on the audio component's rear panel; RCA plugs are the termination of unbalanced interconnects.

XLR Plug and Jack: XLR plugs are three-pin connectors terminating a balanced interconnect. XLR jacks are rear-panel-mounted connectors that accept XLR plugs.

Binding Post: Binding posts are terminations on power amplifiers and loudspeakers that provide connection points for loudspeaker cables.

Five-way Binding Post: A type of binding post that can accept bare wire, spade lugs, pins, round lugs over the post, or banana plugs. Five-way posts are found on most power amplifiers and loudspeakers.

Spade Lug: A flat, pronged termination for loudspeaker cables. Spade lugs fit around power-amplifier and loudspeaker binding posts. The most popular kind of loudspeaker-cable termination.

Banana Plug and Jack: Banana plugs are sometimes found on loudspeaker cables in place of spade lugs. Banana plugs will fit into five-way binding posts or banana jacks. Many European products use banana jacks on power amplifiers for loudspeaker connection.

AWG: American Wire Gauge: a measure of conductor thickness, usually in loudspeaker cables. The lower the AWG number, the thicker the wire. Lamp cord has an AWG of 18, usually referred to as “18 gauge.”

HDMI: An initialism for High-Definition Digital Multimedia Interface, HDMI was developed for home theater to carry high-definition video along with high-resolution digital audio on the same cable. [Chapter 11](#) (“Audio for Home Theater”) includes a detailed description of HDMI.

Note: AC power cords are covered in [Chapter 15](#), along with AC power conditioners.

Some of these cables, interconnects, and terminations are shown in [Fig.10-1](#).



Fig.10-1 From left to right: balanced interconnects terminated with XLR connectors, unbalanced interconnects terminated with RCA plugs, spade lugs, and banana plugs (Courtesy AudioQuest)

How to Choose Cables and Interconnects

Ideally, every component in the system—including cables and interconnects—should be absolutely neutral and impose no sonic signature on the music. As this is never the case, we are forced to select cables and interconnects with colorations that counteract the rest of the system's colorations.

For example, if your system is a little on the bright and analytical side, mellow-sounding interconnects and cables can take the edge off the treble and let you enjoy the music more. If the bass is overpowering and fat, lean- and tight-sounding interconnects and cables can firm up and lean out the bass. A system lacking palpability and presence in the midrange can benefit from a forward-sounding cable.

Selecting cables and interconnects for their musical compatibility should be viewed as the final touch to your system. A furniture maker who has been using saws, planes, and rasps will finish his work with steel wool or very fine sandpaper. Treat cables and interconnects the same way—as the last tweak to nudge your system in the right direction, not as a Band-Aid for poorly chosen components.

Cables and interconnects won't correct fundamental musical or electrical incompatibilities. For example, if you have a high-output-impedance power amplifier driving current-hungry loudspeakers, the bass will probably be soft and the dynamics constricted. Speaker cables won't fix this problem. You might be able to ameliorate the soft bass with the right cable, but it's far better to fix the problem at the source—a better amplifier/loudspeaker match.

Good cables merely allow the system's components to perform at their highest level; they won't make a poor system or bad component match sound good. Start with a high-quality, well-chosen system, and select cables and interconnects that allow that system to achieve its highest musical performance. Remember, a cable or interconnect can't actually effect an absolute improvement in the sound; the good ones merely do less harm.

A typical hi-fi system will need one pair of loudspeaker cables (two pairs for bi-wiring), one pair of long interconnects between the preamplifier and power amplifier, and several pairs of short interconnects for connections between source components (such as a turntable or DAC) and the preamplifier.

If the power amplifier is located near the speakers, the speaker cables will be short and the interconnects between the preamplifier and power amplifier will be

long. Conversely, if the power amplifier is near the source components and preamplifier, the interconnects will be short and the speaker cables long. There is no consensus among the experts about which method is preferable, but I use long interconnects and short loudspeaker cables. Ideally, interconnects *and* loudspeaker cables should be short, but that often isn't practical.

Once you've got a feel for how your system is—or will be—configured, make a list of the interconnects and cables you'll need, and their lengths. Keep all lengths as short as possible, but allow some flexibility for moving loudspeakers, putting your preamp in a different space in the rack, or other possible changes. Although we want to keep the cables and interconnects short for the best sound, there's nothing worse than having interconnects 6" too short. After you've found the minimum length, add half a meter for flexibility.

Interconnects are often made in standard lengths of 1, 1.5, and 2 meters. These are long enough for source-to-preamp connections, but too short for many preamp-to-power-amp runs. These long runs are usually custom-made to a specific length. Similarly, speaker cables are typically supplied in 8' or 10' pairs, but custom lengths are readily available. It's better to have the cable manufacturer terminate the cables (put spade lugs or banana plugs on loudspeaker cables, and RCA or XLR plugs on interconnects) rather than to try to do it yourself.

Concentrate your cable budget on the cables that matter most. The priority should be given to the sources you listen to most often. For example, you may not care as much about the sound of your tuner as you do your DAC. Consequently, you should spend more on interconnects between the DAC and preamplifier than between the tuner and preamp. And because all your sources are connected to the power amplifier through the interconnect between the preamplifier and power amplifier, this link must be given a high priority. But any component—even a mobile device's analog output—will benefit from good interconnects.

Should all of your interconnects and speaker cables be made by the same manufacturer? Or is it better to mix and match brands? There are two schools of thought on this. The first holds that an entire system wired with one brand of cable and interconnect is the best route. If one interconnect works well in your system, use it throughout. This argument assumes that the cable designer made his interconnects and speaker cables to work together to achieve the best possible sound.

The second school of thought holds that different brands are best. Cables and interconnects made by the same manufacturer tend to share a common sonic signature; using the same interconnect and cable throughout the system will only

reinforce that signature. By using cables and interconnects from different manufacturers, the characteristic sonic signature won't be superimposed on the music by every piece of wire.

This second theory has an analog in the recording world. Engineers will record through one brand of recording console, then mix the record through a different brand of console. They don't want to hear the console's sound in the final product, so they don't subject the signal to the same sonic signature twice.

My experience suggests that the only way to determine the best cable or interconnect for your system is to experiment and listen. In some cases, the best results will be achieved with all the interconnects and cables made by the same manufacturer. In others, a mix of different interconnects will work best. It's impossible to predict which cables will sound best in your system.

Most dealers will let you take home several cables at once to try in your system. Take advantage of these offers. Some mail-order companies will send you many cables to try: you keep the ones you want to buy—if any—and return the others. Compare inexpensive cables with expensive ones—sometimes manufacturers have superb cables that sell for a fraction of the price of their top models.

If you're starting a system from scratch, selecting cables is more difficult than replacing one length in your system. Because different combinations of cables will produce different results, the permutations are greatly increased. Moreover, you don't have a baseline reference against which to judge how good or bad a cable is. In this situation, the best way of getting the ideal cables for your system is to follow your dealer's advice. Try the cables and interconnects he suggests, along with two other brands or models for comparison.

How Much Should You Spend on Cables and Interconnects?

At the top end of the scale, cable and interconnect pricing sometimes bears little relationship to the cost of designing and manufacturing the product. Unlike other audio products, whose retail prices are largely determined by the parts cost (the retail price is typically four to six times the cost of raw parts), cables and interconnects are sometimes priced according to what the market will bear. This trend began when one company set its prices vastly higher than everyone else's—and saw its sales skyrocket as a result. Other manufacturers then raised *their* prices so they wouldn't be perceived as being of lower quality. Although some very expensive cables and interconnects are worth the money, many cables are ridiculously overpriced.

The budget-conscious audiophile can, however, take advantage of this phenomenon. Very often, a cable manufacturer's lower-priced products are very nearly as good as its most expensive models. The company prices its top-line products to foster the impression of being "high-end," yet relies on its lower-priced models for the bulk of its sales. When shopping for speaker cables and interconnects, listen to a manufacturer's lower line in your system—even if you have a large cable budget. You may be pleasantly surprised.

Because every system is different, it's impossible to be specific about what percentage of your overall system investment you should spend on cables and interconnects. Spending 5% of your system's cost on cables and interconnects would be an absolute minimum, with about 15% a maximum figure. If you choose the right cables and interconnects, they can be an excellent value. But poor cables on good components will give you poor sound and thus constitute false economy.

Again, I must stress that high cost doesn't guarantee that the cable is good or that it will work well in your system. Don't automatically assume that an expensive cable is better than a low-priced one. Listen to a wide variety of price levels and brands. Your efforts will often be rewarded with exactly the right cable for your system at a reasonable price.

What to Listen For

Cables must be evaluated in the playback system in which they will be used. Not only is the sound of a cable partially system-dependent, but the sonic characteristics of a specific cable will work better musically in some systems than in others. Personal auditioning is the *only* way to evaluate cables and interconnects. Never be swayed by technical jargon about why one cable is better than another. Much of this is pure marketing hype, with little or no relevance to how the cable will perform musically in your system. Trust your ears.

Fortunately, evaluating cables and interconnects is relatively simple; the levels are automatically matched between cables, and you don't have to be concerned about absolute-polarity reversal. One pitfall, however, is that cables and interconnects need time to break in before they sound their best. Before break-in, a cable often sounds bright, hard, fatiguing, congested, and lacking in soundstage depth. These characteristics often disappear after several hours' use, with days or weeks of use required for full break-in. You can't be sure, however, if the cable is inherently bright- and hard-sounding, or if it just needs breaking in. Note that break-in wears off over time. Even if a cable has had significant use, after a long period of not being used it may not sound its best until you've put music through it for a few days.

With those cautions in mind, you're ready to evaluate cables and interconnects. Listen to the first interconnect for 15 minutes to half an hour, then replace it with the next candidate. One way of choosing between them is merely to ask yourself which interconnect allows you to enjoy the music more. You don't need to analyze what you're hearing; just pick the interconnect that makes you feel better.

The other method is to scrutinize what you're hearing from each interconnect and catalog the strengths and weaknesses. You'll often hear trade-offs between interconnects: one may have smoother treble and finer resolution than another, but less soundstage focus and transparency. Another common trade-off is between smoothness and resolution of detail: The smooth cable may lose some musical information, but the high-resolution cable can sound analytical and bright. Again, careful auditioning in your own system is the only way to select the right cables and interconnects. Keep in mind, however, that a better cable can sometimes reveal flaws in the rest of your system. You should also know that

cables and interconnects seem to sound better after they have “settled in” for a few days without being moved.

Cables and interconnects can add some annoying distortions to the music. I’ve listed below the most common sonic problems of cables and interconnects. (A full description of these terms is included in [Chapter 3](#).)

Grainy and hashy treble: Many cables overlay the treble with a coarse texture. The sound is rough rather than smooth and liquid.

Bright and metallic treble: Cymbals sound like bursts of white noise rather than a brass-like shimmer. They also tend to splash across the soundstage rather than sound like compact images. Sibilants (s and sh sounds on vocals) are emphasized, making the treble sound “spitty.” It’s a bad sign if you suddenly notice more sibilance. The opposite condition is a dark and closed-in treble. The cable should sound open, airy, and extended in the treble, without sounding overly bright, etched, or analytical.

Hard textures and lack of liquidity: Listen for a glassy glare on the upper registers of solo piano. Similarly, massed voices can sound glazed and hard rather than liquid and richly textured.

Listening fatigue: A poor cable will quickly cause listening fatigue. The symptoms of listening fatigue are headache, a feeling of relief when the music is turned down or stopped, the need to do something other than listen to music, and the feeling that your ears are tightening up. This last condition is absolutely the worst thing any audio component can do. Good cables (in a good system) will let you listen at higher levels for longer periods of time. If a cable or interconnect causes listening fatigue, avoid it regardless of its other attributes.

Lack of space and depth: Using a recording with lots of natural depth and ambience, listen for how the cable affects soundstage depth and the sense of instruments hanging in three-dimensional space. Cables also influence on the sense of image focus. Poor cables can also make the soundstage less transparent.

Low resolution: Some cables and interconnects sound smooth, but they obscure the music’s fine detail. Listen for low-level information and an instrument’s inner detail. The opposite of smoothness is a cable that’s “ruthlessly revealing” of every detail in the music, but in an unnatural way. Musical detail should be audible, but

not hyped or exaggerated. The cable or interconnect should strike a balance between resolution of information and a sense of ease and smoothness.

Mushy bass or poor pitch definition: A poor-quality cable or interconnect can make the bass slow, mushy, and lacking in pitch definition. With such a cable, the bottom end is soggy and fat rather than taut and articulate. Low-frequency pitches are obscured, making the bass sound like a roar instead of being composed of individual notes.

Constricted dynamics: Listen for the cable or interconnect's ability to portray the music's dynamic structure, on both small and large scales. For example, a guitar string's transient attack should be quick, with a dynamic edge. On a larger scale, orchestral climaxes should be powerful and have a sense of physical impact (if the rest of your system can portray this aspect of music).

I must reiterate that putting a highly colored cable or interconnect in your system to correct a problem in another component (a dark-sounding cable on a bright loudspeaker) isn't the best solution. Instead, use the money you would have spent on new cables on better speakers—*then* go cable shopping. Cables and interconnects shouldn't be Band-Aids; instead, cables should be the finishing touch to let the rest of your components perform at their highest level.

Binding Posts and Cable Terminations

Binding posts vary hugely in quality, from the tiny, spring-loaded, push-in terminal strips on cheap loudspeakers to massive, custom-made, machined brass posts plated with exotic metals. Poor binding posts not only degrade the sound, they also break easily. When shopping for power amplifiers and speakers, take a close look at binding-post quality.

The most popular type is the *five-way binding post*. It accepts spade lugs, banana plugs, pins, round lugs over the post, or bare wire. Some five-ways are nickel-plated; higher-quality ones are plated with gold and won't tarnish. Five-way binding posts should be tightened with a ½" nut driver or tool specifically designed for binding posts, not a socket and ratchet or wrench that could easily overtighten the nut. The connection should be tight, but not to the point of stripping the post or causing it to turn in the chassis. When tightening a five-way binding post, watch the inside ring or collar next to the chassis; if it begins to turn, you've overtightened the post and are in danger of damaging the power amplifier or loudspeaker.

Custom posts of heavy-duty machined metal are more robust than five-way posts—they can accept more torque without stripping or coming loose in the chassis. Custom posts are often found on the most expensive equipment. Some binding posts have such a large center post that spade lugs won't fit around them. These posts have large holes in the center for accepting large bare wires or banana jacks. Although these posts are expensive and appear to be of high quality, they're inconvenient to use. If your equipment has this sort of post, the best solution is to terminate your speaker cables with oversized spade lugs. Most spade lugs have a distance between the prongs of ¼" to 3/16"; oversized spades have prongs 5/16" apart—enough of a difference to fit the large-holed posts.

If you have a choice of bare wire, banana plug, or spade lug on speaker-cable terminations, go with the spade lug. It forms the best contact with a binding post and is the most standard form of connection. Many European products provide only banana jacks, forcing you to use sonically inferior banana-plug terminations.

You should be aware that *any* termination slightly degrades the sound of your system. Consequently, some audiophiles have gone to the trouble of removing all plugs, jacks, spade lugs, and binding posts from their systems and hardwiring everything together. This is an extreme measure and makes switching equipment

difficult or impossible. Hardwiring is an option, but not one that should be undertaken without considerable deliberation and technical expertise.

Bi-Wired Loudspeaker Cables

Bi-wiring is running two lengths of cable between the power amplifier and loudspeaker. This technique usually produces much better sound quality than conventional single-wiring. Most high-end loudspeakers have two pairs of binding posts for bi-wiring, with one pair connected to the crossover's tweeter circuit and the other pair connected to the woofer circuit. The jumpers, installed at the factory that connect the two pairs of binding posts, must be removed for bi-wiring.

In a bi-wired system, the power amplifier “sees” a higher impedance on the tweeter cable at low frequencies, and a lower impedance at high frequencies. The opposite is true in the woofer half of the bi-wired pair. This causes the signal to be split up, with high frequencies traveling mostly in the pair driving the speaker's tweeter circuit and low frequencies conducted by the pair connected to the speaker's woofer circuit. This frequency splitting reportedly reduces magnetic interactions in the cable, resulting in better sound. The large magnetic fields set up around the conductors by low-frequency energy can't affect the transfer of treble energy. No one knows exactly how or why bi-wiring works, but with nearly all loudspeakers with bi-wiring provision, it makes a big improvement in the sound. Whatever your cable budget, you should bi-wire if your loudspeaker has bi-wired inputs, even if it means buying two runs of less expensive cables.

You can bi-wire your loudspeakers with two identical single-wire runs, or with a specially prepared bi-wire set. A bi-wire set has one pair (positive and negative) of terminations at the amplifier end of the cable, and two pairs at the loudspeaker end of the cable. This makes it easier to hook up, and probably offers slightly better sound quality.

Loudspeakers can also be connected with a *single bi-wire* set, in which a single cable with multiple internal conductors has two pairs of terminations on one end and a single pair of terminations at the other end. Although this approach is much less expensive than two runs of cable, you lose the benefit of magnetically isolating the low- and high-frequency conductors from each other. Two bi-wiring options are shown in [Fig.10-2](#).

Most bi-wired sets use identical cables for the high- and low-frequency legs. Mixing cables, however, can have several advantages. By using a cable with good bass on the low-frequency pair, and a more expensive but sweeter-sounding cable on the high-frequency pair, you can get better performance for a lower cost. Use a

less expensive cable on the bass and put more money into the high-frequency cable. If you've already got two pairs of cable the same length, the higher-quality cable usually sounds better on the high-frequency side of the bi-wired pair. If you use different cables for bi-wiring, they should be made by the same manufacturer and have similar physical construction. If the cables in a bi-wired set have different capacitances or inductances, those capacitances and inductances change the speaker's crossover characteristics.




Fig.10-2 Loudspeakers can be connected with one cable (top), a single bi-wire connection (middle), with a double-bi-wire cable (bottom), or with two completely separate runs of cable (not shown). (Courtesy AudioQuest)

Balanced and Unbalanced Lines

Line-level interconnects come in two varieties: *balanced* and *unbalanced*. A balanced interconnect is recognizable by its three-pin XLR connector. An unbalanced interconnect is usually terminated with an RCA plug. Balanced and unbalanced lines are shown in the photograph at the beginning of this chapter.

Why do we use two incompatible systems for connecting equipment? At one time, all consumer audio hardware had unbalanced inputs and outputs, and all professional gear was balanced. In fact, balanced inputs are often called “professional inputs,” to differentiate them from “consumer” unbalanced jacks. Balanced connection was considered both unnecessary and too expensive for home playback systems.

The emergence of high-end audio changed that thinking. Instead of using the least expensive connection method, high-end product designers started using higher-quality balanced lines and terminations in consumer gear. The better the equipment, the more likely that it has at least some balanced connections in addition to unbalanced jacks. Moreover, more and more manufacturers are offering balanced connections on their equipment. This is why we have two standards—balanced and unbalanced. The technical and sonic advantages of balanced connection, once the exclusive domain of audio professionals, are becoming increasingly available for home playback systems.


But what exactly is a balanced line, and how is it different from a standard RCA cable and jack?

In an unbalanced line, the audio signal appears across the center pin of the RCA jack and the shield, or ground wire. Some unbalanced interconnects have two signal conductors and a shield, with the shield not used as a signal conductor. If this unbalanced interconnect happens to be close to fluctuating magnetic fields—an AC power cord, for example—the magnetic field will induce a noise signal in the interconnect that is heard as hum and noise reproduced through the playback system’s speakers.


In a professional application, this hum, buzz, and noise are unacceptable, leading to the development of an interconnection method that is immune to noise interference: the balanced line. A balanced line has three conductors: two carrying signal, and one ground. The two signals in a balanced line are identical, but 180° out of phase with each other. When the signal in one of the conductors is at peak positive, the signal in the other conductor is at peak negative (see

Fig.10-3). The third conductor is signal ground. Some balanced interconnects use three conductors plus a shield.

When the two identical but opposite polarity signals carried on the balanced line are input to a *differential amplifier* in the component receiving the signal, noise picked up by the interconnect is rejected. Here's why: a differential amplifier amplifies only the *difference* between the two signals (see Fig.10-4). If noise is introduced into the line, the noise will be common to both conductors and the differential amplifier will reject the noise. This phenomenon of rejecting noise signals common to both conductors in a balanced line is called *common-mode rejection*. Differential inputs are specified according to how well they reject signals common to both conductors, a measurement called *common-mode rejection ratio*, or *CMRR*. Note that a balanced line won't make a noisy signal clean; it just prevents additional noise from being introduced in the interconnect. If the noise is common to both halves of the balanced line, however, common-mode rejection will eliminate the noise.




(A) Unbalanced line (using two conductors)




(B) Balanced line (using three conductors)

Fig.10-3 Signal voltages applied to an unbalanced and a balanced line (Copied with permission from *Audio Technology Fundamentals*, by Alan A. Cohen)




(A) Required signal

(B) Noise interference



(C) Unbalanced signal and noise amplification



(D) Balanced line—signal amplification and noise elimination

Fig.10-4 When the two signals carried by the balanced line are input to a differential amplifier, noise common to both conductors is rejected (bottom). (Copied with permission from *Audio Technology Fundamentals*, by Alan A. Cohen)

In professional applications, a transformer sometimes serves the same function as the differential amplifier, passing only the difference between the two signal conductors and rejecting signals common to both conductors. Consumer audio equipment uses differential amplifiers, not transformers, but the concept is identical.

In a balanced line terminated with XLR connectors, pin 1 is always signal ground. There is, however, no universal convention for which of the two signal conductors carries the non-inverted signal and which carries the inverted signal. The non-inverted conductor is often called the “hot” conductor, with the inverted conductor designated “cold.” After decades of no clear standard, the Audio

Engineering Society recently adopted the North American tradition of having pin 2 carry the non-inverted (“hot”) signal, pin 3 the inverted (“cold”) signal.

How the balanced line is wired (pin 2 or pin 3 “hot”) can determine if your playback system is inverting or not inverting absolute polarity (see the section on absolute polarity in [Appendix A](#)). If your system is non-inverting—that is, a positive-going signal on the recording produces a positive-going signal at the speakers—substituting a power amp that has balanced inputs with pin 3 “hot” with one with pin 2 “hot” will make your system inverting. When switching balanced components—DACs, preamps, or power amps—you should know if the new component’s XLR jacks are wired the same—either pin 2 or pin 3 “hot”—as the existing component. You can also change your system’s absolute polarity by rewiring the balanced interconnect to swap the wires going to pins 2 and 3. But if you want to change your system’s absolute polarity, it’s far better to simply reverse absolute polarity by switching red for black and black for red on both speaker cables. The freak interconnect with pins 2 and 3 switched at one end may wind up in another system or application where you *don’t* want the polarity switched.

Quite apart from the advantage of noise cancellation in a balanced line, a balanced connection often sounds better than an unbalanced connection. A system connected by balanced interconnects can, however, also often sound *less* good than one connected with unbalanced lines. Say you have a DAC that takes an unbalanced signal from the digital-to-analog converter chip and converts it to a balanced signal so that the DAC manufacturer can tout the product as having “balanced outputs” (see the section on balanced DACs in [Chapter 7](#)). Inside the DAC, the unbalanced signal is converted to a balanced signal by a *phase splitter*, a circuit that takes a signal of one polarity and turns it into two signals of opposite polarity. Phase splitting subjects the unbalanced signal to an additional active (transistor- or op-amp-based) stage and puts more circuitry in the signal path.

The balanced DAC’s output is then input to a balanced-input preamp. Because all but the very best balanced preamplifiers convert a balanced input signal to an unbalanced signal for the preamp’s internal gain stages, the preamp’s input converts this balanced signal to an unbalanced signal—adding yet another active stage to the signal path. After the unbalanced signal is amplified within the preamp, it is converted back to balanced with another phase splitter.

The preamplifier’s balanced output is then sent from the preamplifier output to the power amplifier’s balanced input, where it’s—that’s right—converted to unbalanced by yet another active stage. The result of these unbalanced/balanced/unbalanced/balanced/unbalanced conversions is additional

electronics in the signal path—just what we don't want. This is why you can't assume that balanced components sound inherently better than unbalanced ones. Magazine reviews of audio components should include musical *and* technical comparisons of the product's balanced and unbalanced modes.

Some products, however, are truly balanced and don't rely on phase splitters and unbalancing amplifiers. For example, a DAC may create a balanced signal in the digital domain (at no sonic penalty but, indeed, a sonic gain) and convert that balanced signal to analog with four digital-to-analog converters and analog output stages (left channel + and -, right channel + and -). Similarly, some preamplifiers are truly balanced and have double the circuitry to operate on the non-inverting and inverting signals separately. You can tell a truly balanced preamplifier by the number of elements in its volume control. A preamplifier that internally operates on an unbalanced signal will have two volume-control elements: one for the left channel, one for the right. A fully balanced preamplifier will have four elements: \pm left channel and \pm right channel. The signal thus stays balanced from before the DACs inside the digital processor all the way to the final stages in the power amplifier.

As in all things audio, the proof is in the listening. When shopping for a component, listen to it in both balanced and unbalanced modes. Let your ears decide if the component works best in your system when connected via the balanced or unbalanced lines.

Cable and Interconnect Construction

Cables and interconnects are composed of three main elements: the signal conductors, the dielectric, and the terminations. The *conductors* carry the audio signal; the *dielectric* is an insulating material between and around the conductors; and the *terminations* provide connection to audio equipment. These elements are formed into a physical structure called the cable's *geometry*. Each of these elements—particularly geometry—can affect the cable's sonic characteristics.

Conductors

Conductors are usually made of copper or silver wire. In high-end cables, the copper's purity is important. Copper is sometimes specified as containing some percentage of "pure" copper, with the rest impurities. For example, a certain copper may be 99.997% pure, meaning it has three-thousandths of one percent impurities. These impurities are usually iron, sulfur, antimony, aluminum, and arsenic. Higher-purity copper—99.99997% pure—is called "six nines" copper. Many believe that the purer the copper, the better the sound. Some copper is referred to as *oxygen-free copper (OFC)*. This is copper from which the oxygen molecules have been removed. It is more proper to call this "oxygen-reduced" copper, because it is impossible to remove all the oxygen. In practice, OFC has about 50 parts per million (ppm) of oxygen compared to 250ppm of oxygen for normal copper. Reducing the oxygen content retards the formation of copper oxides in the conductor, which can interrupt the copper's physical structure and degrade sound quality.

Another term associated with copper is *linear crystal (LC)*, which describes the copper's structure. Drawn copper has a grain structure that can be thought of as tiny discontinuities in the copper. The signal can be adversely affected by traversing these grains; the grain boundary can act as a tiny circuit, with capacitance, inductance, and a diode effect. Standard copper has about 1500 grains per foot; LC copper has about 70 grains per foot. Fig.10-5 shows the grain structure in copper having 400 grains per foot. Note that the copper isn't isotropic; it looks decidedly different in one direction than the other. All copper made into thin wires exhibits a chevron structure, shown in the photograph in Fig.10-5. This chevron structure may explain why some cables sound different when reversed.

Conductors are made by casting a thick rod, then drawing the copper into a smaller gauge. Another technique—which is rare and expensive—is called "as-cast." This method casts the copper into the final size without the need for drawing.

The highest-quality technique for drawing copper is called Ohno Continuous Casting or OCC. OCC copper has one grain in about 700 feet—far less than even LC copper. The audio signal travels through a continuous conductor instead of traversing grain boundaries. Because OCC is a process that can be performed on any purity of copper, not all OCC copper is equal.



Fig.10-5 Copper in cables and interconnects has a grain structure. (Courtesy AudioQuest)

The other primary—but less prevalent—conductor material is silver. Silver cables and interconnects are obviously much more expensive to manufacture than copper ones, but silver has some advantages. Although silver's conductivity is only slightly higher than that of copper, silver oxides are less of a problem for audio signals than are copper oxides. Silver conductors are made using the same drawing techniques used in making copper conductors.

The Dielectric

The dielectric is the material surrounding the conductors, and is what gives cables and interconnects some of their bulk. The dielectric material has a large effect on the cable's sound; comparisons of identical conductors and geometry, but with different dielectric materials, demonstrate the dielectric's importance.

Dielectric materials absorb energy, a phenomenon called *dielectric absorption*. A capacitor works in the same way: a dielectric material between two charged plates stores energy. But in a cable, dielectric absorption can degrade the signal. The energy absorbed by the dielectric is released back into the cable slightly delayed in time—an undesirable condition.

Dielectric materials are chosen to minimize dielectric absorption. Less expensive cables and interconnects use plastic or PVC for the dielectric. Better cables use polyethylene; the best cables are made with polypropylene or even Teflon dielectric. One manufacturer has developed a fibrous material that is mostly air (the best dielectric of all, except for a vacuum) to insulate the conductors within a cable. Other manufacturers inject air in the dielectric to create a foam with high air content. Just as different dielectric materials in capacitors sound different, so too do dielectrics in cables and interconnects.

Terminations

The terminations at the ends of cables and interconnects are part of the transmission path. High-quality terminations are essential to a good-sounding cable. We want a large surface contact between the cable's plug and the component's jack, and high contact pressure between them. RCA plugs will sometimes have a slit in the center pin to improve contact with the jack. This slit is effective only if the slit end of the plug is large enough to be compressed by insertion in the jack. Most high-quality RCA plugs are copper, with some brass mixed in to add rigidity. This alloy is plated with nickel, then flashed with gold to prevent oxidation. On some plugs, gold is plated directly on the brass. Other materials for RCA plugs and plating include silver and rhodium.

RCA plugs and speaker-cable terminations are soldered or welded to the conductors. Most manufacturers use solder with some silver content. Although solder is a poor conductor, the spade lugs are often crimped to the cable first, forming a "cold" weld that makes a gas-tight seal. In the best welding technique, resistance welding, a large current is pulsed through the point where the conductor meets the plug. The resistance causes a small spot to heat, melting the two metals. The melted metals merge into an alloy at the contact point, ensuring good signal transfer. No matter what the termination method, a mechanism inside the cable, called a strain relief, protects the termination from physical stress.

Geometry

How all of these elements are arranged constitutes the cable's geometry. Some designers maintain that geometry is the most important factor in cable design—even more important than the conductor material and type.

A cable's performance: can be affected by something as simple as twisting a pair of conductors around each other instead of running them side by side. Twisting the conductors greatly reduces capacitance and inductance in the cable. Think of the physical structure of two conductors running in parallel, and compare that to the schematic symbol for a capacitor, which is two parallel lines.

This is the grossest example; there are many fine points to cable design. I'll describe some of them here, with the understanding that I'm merely presenting certain opinions on cable construction, not endorsing a particular method.

Most designers agree that the *skin effect*, and interaction between strands, are the greatest sources of sonic degradation in cables. In a cable with high skin effect, more high-frequency signal flows along the conductor's surface, less through the conductor's center. This occurs in both solid-core and stranded conductors (Fig.10-6). Skin effect changes the cable's characteristics at different depths, causing different frequency ranges of the audio signal to be affected by the cable differently. The musical consequences of skin effect include loss of detail, reduced top-octave air, and truncated soundstage depth.




Fig.10-6 Skin effect causes more of the audio signal to travel on the outside of the conductor. (Courtesy AudioQuest)

A technique for battling skin effect is *litz* construction, which simply means that each strand in a bundle is coated with an insulating material to prevent it

from electrically contacting the strands around it. Each small strand within a litz arrangement will have virtually identical electrical properties. Litz strands push skin-effect problems out of the audible range. Because litz strands are so small, many of them bundled together in a random arrangement are required to achieve a sufficient gauge to keep the resistance low.

A problem with stranded cable (if it isn't of litz construction) is a tendency for the signal to jump from strand to strand if the cable is twisted. One strand may be at the outside at a point in the cable, then be on the inside farther down the cable. Because of skin effect, the signal tends to stay toward the outside of the conductor, causing it to traverse strands. Each strand interface acts like a small circuit, with capacitance and a diode effect, much like the grain structure within copper.

Individual strands within a conductor bundle can also interact magnetically. Whenever current flows down a conductor, a magnetic field is set up around that conductor. This alternating magnetic field can induce a signal in adjacent conductors (see [Appendix B](#)), and degrade the sound. Some cable geometries reduce magnetic interaction between strands by arranging them around a center dielectric, which keeps them farther apart. These are just a few of the techniques used by cable designers to make better-sounding cables.

Terminated Cables and Interconnects

Some cables and interconnects are not merely pieces of wire, but contain electronic components. These cables are identified by “boxes” on one or both cable ends housing the resistors, inductors, and capacitors that form an electrical “network.” These products are often called terminated cables. To my knowledge, two companies offer terminated cables and interconnects: Music Interface Technologies (MIT) and Transparent Audio (Figs.10-7 and 10-8).

According to MIT, some of the audio signal’s voltage is stored in the cable’s capacitance, and some of the signal’s current is stored in the cable’s inductance. The amount of energy stored in the cable varies with frequency, which is manifested as tonal and dynamic emphases at those frequencies. Moreover, the energy stored in the cable is released over time, rather than being delivered to the loudspeaker along with the rest of the signal. This frequency-dependent, non-linear energy storage distorts the size and shape of the soundstage, according to MIT.

Terminated cables are designed to prevent such energy storage, and to deliver all the signal to the loudspeaker with the correct phase (timing) relationship. MIT claims that terminated cables are essential to realizing correct bass weight, smooth tonal balance, and a full-sized soundstage with precise image focus.

Because terminated cables form a low-pass filter (i.e., they don’t start filtering until about 1MHz, well above the audioband), they are useful for connecting very-wide-bandwidth electronics. The Spectral brand of electronics, for example, has bandwidths as high as 3MHz—not because the designer wants to pass those frequencies, but to make the circuits behave better in the audioband. Terminated cables filter those very high frequencies, ensuring that loudspeakers are never fed energy in the megahertz range, and making the system more stable.



Fig.10-7 Terminated loudspeaker cables (Courtesy Music Interface Technologies)




Fig.10-8 Terminated interconnects (Courtesy Transparent Audio)

Battery Bias in Cables and Interconnects

Anyone who has listened to speaker cables and line-level interconnects will have discovered that cables and interconnects start to sound better (smoother treble, greater ease, less grain, more dimensional soundstage) after they've been in the system for a few days. This phenomenon is often called break-in. It's easy to understand break-in of speakers and phono cartridges (the rubber surround in the loudspeaker drivers and the rubber suspension in cartridges relaxes with use), but why would putting audio signals down a conductor cause that conductor's sound to change?

The answer is that the cable's dielectric "forms" to the charged state created by current flowing through the conductor. The same phenomenon occurs in capacitors, and is one reason electronics sound better when left on for a few days. After a cable or capacitor "forms," it returns to its uncharged state slowly. This means that cables that have been broken in will slowly return to their original state when the system is not in use. The process is not fully understood, but extensive listening experience confirms this cycle.

Rather than rely on the audio signal to "form" the cable's dielectric, the cable designer can apply a relatively high DC voltage across the dielectric when the cable is made, in essence breaking-in the cable more effectively than could an audio signal, and keeping the cable in that state indefinitely.

That is the theory behind a new breed of cable that is supplied with a battery attached. AudioQuest, which has applied for a patent on the idea, calls the technique Dielectric-Bias System (DBS). A wire connected to a 24V battery's positive terminal runs down the cable's length, and is not connected to the cable or to the battery's negative terminal. The battery's negative terminal is connected to a shield around the conductors. Because there is no path for current between the battery's positive and negative terminals, the battery experiences no current drain and will last for years. Keep in mind that the battery system is completely outside the audio signal path, and doesn't interact with the audio signal in any way. An interconnect employing this technology is shown in [Fig.10-9](#).




Fig.10-9 A loudspeaker cable employing the Dielectric-Bias System (Courtesy AudioQuest)

A related technique is employed by Richard Vandersteen in his Model 5 and Model 7 loudspeakers. The two speaker models use an integral powered woofer and an external crossover to filter bass from the signal driving the speaker. Vandersteen included a battery-bias system across the capacitors in this crossover to keep the capacitors biased at a relatively high voltage. This DC shifts the audio signal away from the zero-crossing point, improving sound quality. This technique is different from AudioQuest's Dielectric-Bias System in that the audio signal is combined with the battery's DC voltage.

Cable and Interconnect Specifications

There's a lot of hype and just plain misinformation about cables and interconnects. Manufacturers sometimes feel the need to invent technical reasons for why their cables sound better than the competition's. In reality, cable design is more art than science, with good designs emerging from trial and error (and careful listening). Although certain conductors, dielectrics, and geometries have specific sonic signatures, successful cable designs can seldom be described in technical terms. This is why cables should never be chosen on the basis of technical descriptions and specifications.

Nonetheless, some cable and interconnect specifications should be considered in some circumstances. The three relevant specifications are *resistance*, *inductance*, and *capacitance*. (These terms are explained in detail in [Appendix B](#).)

A cable or interconnect's resistance, more properly called *DC series resistance*, is a measure of how much it opposes the flow of current through it. The unit of resistance is the *ohm*. The lower the number of ohms, the lower the cable's or interconnect's resistance to current flow. In practice, cable resistance is measured in tenths of ohms. Resistance isn't usually a factor in interconnect performance (except in some of the new nonmetallic types), but can affect some speaker cables—particularly thin ones—because of their higher current-carrying requirements.

The sounds of interconnects and speaker cables can be affected by inductance. It is generally thought that the lower the inductance, the better, particularly in speaker cables. Some power amplifiers, however, need to see some inductance to keep them stable; many have an output inductor connected to the speaker binding posts (inside the case). When considering how much inductance the power amplifier sees, you must add the cable inductance to the speaker's inductance.

Capacitance is an important characteristic of interconnects, particularly when long runs are used, or if the source component has a high output impedance. Interconnect capacitance is specified in picofarads (pF) per foot. What's important isn't the interconnect's intrinsic capacitance, but the total capacitance attached to the source component. For example, 5' of 500pF-per-foot interconnect has the same capacitance as 50' of 50pF-per-foot interconnect. High interconnect capacitance can cause treble rolloff and restricted dynamics. (A full technical discussion of interconnect capacitance is included in [Appendix B](#).)

Cables in the Power Amplifier/Loudspeaker Interface

The interface between a power amplifier and a loudspeaker through a cable is a critical point in a playback system. Unlike interconnects, which carry low-level signals, speaker cables carry much higher voltages and currents. Speaker cables thus react more with the components they are connected to.

Damping factor is an amplifier's ability to control the woofer's motion after the drive signal has ceased. For example, if you drive a loudspeaker with a bass-drum whack, the woofer's inertia and resonance in the enclosure will cause it to continue moving after the signal has died away. This is a form of distortion that alters the music signal's dynamic properties. Fortunately, the power amplifier can control the motion; the degree of this control, or damping factor, is expressed as a simple number.

Damping factor is related to the amplifier's output impedance. The lower the output impedance, the higher the damping factor. When you connect a power amplifier and loudspeaker with cable, the cable's resistance decreases the amplifier's effective damping factor. For example, an amplifier's damping factor of 100 may be reduced to 40 by 20' of moderately resistive speaker cable. The result is reduced tightness and control in the bass. Speaker cables should therefore have low resistance and be as short as possible.

11 Audio for Home Theater

Introduction

The massive growth of home theater over the past 15 years has had a profound effect on high-end audio. On one hand, home theater—the combination of high-quality audio and video presentation in the home—has opened new markets to high-end companies and new possibilities for the music enthusiast. On the other hand, combining a music-reproduction system with a video display and multichannel surround sound all too often compromises the system’s musical performance. As audiophiles, we go to great lengths to achieve the highest possible sound quality and get closer to the music. Any compromise of that pursuit is anathema to the goals and spirit of high-performance audio.

Yet the application of high-end products to movie watching can create an enormously compelling experience. Film is as important an art form as music; just as good technical quality in music reproduction brings us closer to the musicians’ expression, home theater done right conveys to us more of the filmmakers’ art. Moreover, the availability of a significant number of multichannel titles in the SACD format has fueled a push toward more than two loudspeakers in the listening room—loudspeakers that can reproduce film soundtracks as well as multichannel music programs.

I explored this conflict between high-end audio and home theater in the following short essay, first published as “From the Editor,” in the March, 2003 issue of *The Absolute Sound*.

“Now I know why I hate home theater,” said a colleague to me after a series of home-theater demonstrations at a Consumer Electronics Show. The fellow journalist, a dedicated 2-channel music enthusiast, voiced a familiar refrain in high-end audio circles—namely, that music is a pure art form and home theater is

merely glorified television at best, and an abomination at worst. At the time (1992), I had to agree; the demonstration consisted of subjecting a group of audio writers to a series of the loudest and most violent five-minute clips from the loudest and most violent movies at ear-splitting levels. The manufacturer was trying to impress us, but succeeded only in perpetuating the myth that home theater and high-quality music reproduction are antithetically opposed.

“I say “myth” because since that demonstration ten years ago, I’ve come to believe that the antipathy between music and home theater is a false dichotomy. High-end audio equipment is about getting closer to the artistic expression of composers and musicians. Home theater is (or should be) about getting closer to the artistic expression of filmmakers. The greater the technical quality of the presentation, the more deeply we can experience the artists’ intent—whether musical or cinematic. So what went wrong with this picture?

“The home-theater industry has, unfortunately, given those of us in the high end many reasons to abhor its intrusion into music reproduction: gimmicky surround-sound modes; amusical, pant-rattling subwoofer bass; excessive volume levels; and an aesthetic that values impact over nuance. But an entire concept—the pursuit of high-quality film reproduction in the home—cannot be dismissed solely because of the excesses of those who appeal to the lowest common denominator. Rather than viewing home theater as a bastardization of audio, I’ve come to the belief that home theater, when done right, is an expansion of the high-end ethic, and one that forges an entirely new path for creating an immensely rewarding connection between artist and audience.

“This expansion has been made possible by the entry of true high-end audio companies into the home-theater arena. These companies, driven by high-end ideals, create products that are faithful to the source, eschew gimmicks, and don’t sacrifice musical performance for explosions and sound effects.

“Moreover, home theater has been a boon to high-end audio in several ways. First, the exploding market for home-theater equipment has kept many traditional 2-channel companies prosperous. This allows them to stay in business, and continue to offer 2-channel as well as multichannel products. Second, home theater causes many who once had a “stereo” with loudspeakers in opposite corners of a room to set up the loudspeakers correctly and, for the first time, consider the spatial aspects of reproduced music. Third, a decent-quality home theater that’s well set up can inadvertently become a vehicle for rediscovering the joys of music. I once recommended the Paradigm Atom multichannel loudspeaker system to an “average joe” friend (\$1106 complete at the time). I ran into him a few weeks later, and the first thing he said to me was that he and his

wife now spend much more time listening to music, and that they are happily rediscovering their music collection. Multiply this scenario by millions, factor in upward mobility, and you have a recipe not only for a robust high-end audio industry, but more importantly, the return of music to a prominent place in daily life.”

With that introduction, let’s look at the nuts and bolts of high-quality audio for home theater. Note that this chapter cannot cover the field in the same detail as other topics throughout this book, and I’ve avoided describing the video aspects of home theater.

Overview of Home-Theater Systems

Home theater is made possible by multichannel surround sound. Instead of two speakers in front of you, as in a stereo system, home theater surrounds you with five, six, or even seven. Three speakers are placed in the front of the room, with two, three, or four surround speakers located to the rear and/or sides of the listening/viewing position. Surround sound is what provides the sense of envelopment that makes the home-theater experience so engrossing.

Multichannel surround sound is encoded on DVDs, Blu-ray Discs, satellite transmissions, and cable feeds in the Dolby Digital format. (Blu-ray Discs can include more advanced surround-sound formats, covered later in this chapter.) The Dolby Digital format, introduced in 1997 to replace Dolby Surround (and its decoding variant, Dolby Pro Logic), encodes six discrete audio channels in a single digital bitstream. The six channels are called left, center, right, left surround, right surround, and low-frequency effects (LFE), and correspond to loudspeakers arrayed in the pattern shown in [Fig.11-1](#). The LFE channel carries only low-bass information (20–100Hz), and can be reproduced by any of the loudspeakers, or by a separate subwoofer. Typically, the LFE channel drives a subwoofer. With five full-bandwidth channels plus the 100Hz LFE channel, we call Dolby Digital a 5.1-channel format. The “.1” in 5.1-channel sound is the LFE channel. The competing Digital Theater Systems (DTS) surround-sound format, found on some DVDs and Blu-ray Discs as an option, also delivers 5.1 discrete audio channels to your home-theater system.




Fig.11-1 A home-theater loudspeaker array uses three front speakers, two surround speakers, and a subwoofer. (Courtesy DTS)

The newer, more sophisticated surround-sound formats available on Blu-ray Disc (Dolby Digital Plus, DTS HD, Dolby TrueHD, and DTS HD Master Audio) provide for seven discrete channels plus a subwoofer (7.1-channel sound). We'll look at these formats in detail later in this chapter.

The job of decoding the Dolby or DTS bitstreams into separate audio channels, and then converting those digital data to analog signals, falls to the home-theater controller, also known as a surround-sound processor. In addition to performing surround-sound decoding, the controller receives audio and video signals from the source components (satellite, disc player, tuner, music server) and selects which one is decoded and amplified by your home-theater audio system and sent to your video monitor for display. The controller also performs digital signal processing, adjusts the overall volume, and fine-tunes the levels for individual channels.

The controller's output is eight line-level signals: the left, center, right, left surround, right surround, left back, right back, and subwoofer channels. Note that a home-theater system can be based on a 5.1-channel or a 7.1-channel speaker array. A 7.1-channel array uses four surround speakers, two on the sidewalls and two on the rear wall. The newer surround formats—Dolby Digital Plus, Dolby TrueHD, DTS HD, and DTS HD Master Audio, are all 7.1-channel

formats, though they can be reproduced through a 5.1-channel system. All modern controllers have eight audio channels so that they are compatible with the 7.1-channel surround-sound formats.

The controller's eight separate outputs feed a 7-channel power amplifier (plus an optional subwoofer), where they are amplified to a level sufficient to drive the home theater's loudspeaker system. A home-theater power amplifier can have five, six, or seven amplifier channels in a single housing. Alternately, some power amps have three channels, others have two, and some are available as monoblocks: one amplifier channel per housing. A 3-channel power amplifier allows you to use your existing 2-channel stereo power amplifier on the surround channels and the 3-channel amplifier on the front three loudspeakers, for example. Most subwoofers have built-in power amplifiers to drive their woofer cones.

An audio/video receiver (AVR) incorporates, in one housing, an A/V controller and multichannel power amplifier. AVRs are less expensive than separate controllers and power amplifiers. If you're installing a second system dedicated to home theater, the AVR is a good option. If you're adding home-theater capability to an existing 2-channel music system, or building a dual-purpose system from scratch, a separate controller and power amplifier is the better route.

A home-theater system's left and right speakers reproduce mostly music and sound effects. The center speaker's main job is reproducing dialogue, and anchoring onscreen sound effects on the television screen. Having three speakers across the front of the room helps the sound's location to more closely match the location of the sound source in the picture. For example, in a properly set-up home-theater system, if a car crosses from the left side of the picture to the right, you hear the sound of the car move from the left speaker, through the center speaker, and then to the right speaker. The sound source appears to follow the image on the screen.

The surround speakers have a different job. They're generally smaller than the front speakers, and handle much less energy. Consequently, they can be mounted unobtrusively on or inside a wall. Surround speakers mostly reproduce "atmospheric" or ambient sounds, creating a diffuse aural atmosphere around the listener. In a jungle scene, for example, the surround speakers would re-create sounds such as chirping birds, falling raindrops, and blowing wind; in a city scene, the viewer would be surrounded by traffic sounds. The surround speakers' contribution is subtle, but vitally important to the overall experience. Correctly

set-up surround speakers should not be able to be heard directly, but should instead envelop the viewer/listener in a diffuse soundfield.

Should You Choose a 5.1-, 7.1-, or 9.1-Channel System?

A fundamental question facing the home-theater shopper is whether to buy a 5.1-, 7.1-, or 9.1-channel audio system. The vast majority of today's film soundtracks are encoded in 5.1 channels, but newer films are mixed in 7.1 channels for theatrical exhibition as well as for home theater. The newer surround-sound formats of Dolby Digital Plus, Dolby TrueHD, DTS HD, and DTS HD Master Audio are all based on 7.1 channels (although they can be played through a 5.1-channel speaker array). In a 9.1-channel system with Dolby Pro Logic IIz decoding, the two additional channels drive the “height” speakers positioned in the front of the room, above the left and right speakers.

There are several arguments for choosing a 5.1-channel system. First, most films are mixed in 5.1 channels (many older films are monophonic or 2-channel). Second, mounting and running cables to additional surround-back speakers isn't worth the small improvement in performance (except in very large rooms). Third, many extremely satisfying home-theater systems are based on 5.1-channel playback. Fourth, a 7.1-channel loudspeaker package distributes your speaker expenditure over seven speakers plus a subwoofer rather than over five speakers and the sub, which could compromise the quality—two good surround speakers are better than four mediocre ones. Finally, if your living room dictates that your couch be positioned against the wall opposite the video display, there's little point in adding surround-back speakers on that wall; it's not likely that they'll perform as intended.

The argument for purchasing a 7.1-channel system goes like this: Even though most films are mixed in 5.1 channels, today's A/V receivers and controllers have processing circuits to synthesize the surround-back channels from 5.1-channel sources. Although this trick doesn't provide true 7.1-channel playback (the surround-back signals are simply derived from the left- and right-surround channels), it nonetheless increases the feeling of immersion in the soundfield. With true 7.1-channel sources and systems, the film-sound mixers are able to position sounds directly behind you—an impossible feat with just 5.1 channels. Moreover, we are moving into an increasingly 7.1-channel world with the audio formats on Blu-ray Disc; it's better to get a 7.1-channel system now rather than have to upgrade later.

There's another reason to buy a 7.1-channel system; the Dolby Pro Logic IIz format, which adds two “height” speakers in the front of the room, requires a 7.1-

channel system at a minimum. It is possible to use the surround-back channels in a 7.1-channel system to drive the front height channels. This arrangement provides the height-channel reproduction without the expense of a 9.1-channel system. The thinking is that the two additional channels (over and above the traditional five channels) are better used to reproduce the height channel than to reproduce the surround-back channels.

Here's my view: If you have a moderately sized room and you're adding surround speakers for the first time, go with a 5.1-channel system. If your room is large, or if you're building a theater room as part of new construction, select a 7.1-channel package. Larger rooms benefit from surround-back speakers because they help fill the "hole" between widely spaced left and right surround speakers. If you're building a home with surround-speaker wiring in the walls, there's very little cost and effort required to run wires for surround-back speakers.

Let's take a closer look at the components of a home-theater system and how to choose the features best suited to your system.

Home-Theater Controllers

No product better exemplifies the fundamental shift in home-entertainment technology than the *controller*. The controller combines many diverse functions and sophisticated technologies in a single housing. A modern controller (also called a preamp/processor), shown in Fig.11-2, performs the source-switching functions of a preamplifier, decodes all the surround-sound formats, provides eight channels of digital-to-analog conversion, and incorporates an electronic crossover to split up the audioband. While power amplifiers and speakers change relatively little over time, the controller has undergone a radical transformation in recent years to keep up with new surround-sound formats, digital signal processing, and connectivity.



Fig.11-2 A modern high-end digital controller. (Courtesy Classé Audio)

Specifically, a controller is used to perform these major functions:

1) receive video and audio signals from various source components (DVR, dish, disc player, etc.) and selects which are sent to the video monitor and home-theater audio system (this function is called “source switching”);

2) perform surround decoding;

3) control playback volume;

4) make adjustments in system setup, such as individual channel levels;

5) direct bass to the appropriate loudspeakers or subwoofers (called bass management); and

6) perform THX processing (if so equipped).

A controller has multiple audio and video inputs to accept signals from your source components. It then performs surround decoding to create a multichannel signal for output to the separate power amplifiers. As mentioned earlier, one of the controller’s outputs is marked LFE (Low-Frequency Effects): the “.1” channel in 5.1-channel and 7.1-channel sound, which carries additional bass signals for high-impact, very low bass sounds, such as explosions. The LFE output is connected to a powered subwoofer. Bass from other channels can be selectively added to form the subwoofer drive signal. For example, if you have small front speakers, you can direct the front-channel bass to the subwoofer along with the LFE channel. The bass that would have driven the front speakers is instead mixed with the LFE signal and sent to the subwoofer, relieving the front speakers from the burden of reproducing low bass.

The controller is the master control center of your home-theater system. When you first set up a controller, its onscreen display will guide you through selecting such things as whether or not you’re using a subwoofer, if your surround loudspeakers are large or small, and calibrating individual channel levels and delay times. You will also control the system on a day-to-day basis through the controller’s front panel and remote control. This includes source selection, volume adjustment, and the setting of individual channel levels (described later in this chapter).

Today’s controllers may also offer a range of features for accessing file-based music, including wireless streaming and network connectivity. These technologies are described in [Chapter 8](#).

How to Choose a Controller:

Controllers range from entry-level units costing about \$1500 to state-of-the-art models that will set you back as much as \$20,000. This wide price range reflects the significant differences between controllers in their features, connectivity, future expandability, and sonic performance.

Let's look at some controller features so that you'll know what to look for when shopping for a controller.

Inputs, Outputs, and Source Switching:

Let's start with the controller's most basic function: selecting the source you listen to or watch. The controller accepts audio or A/V (audio and video) signals from all your source components and lets you select which source signal is sent to the power amplifiers and video monitor. A basic controller will offer two analog-audio inputs (for a tuner and disc player, for example), three digital audio inputs, and perhaps four audio/video (A/V) inputs, with the last on HDMI connectors. In addition to the main outputs that drive your TV and power amplifiers, a record output is often provided to drive a recording device.

When choosing a controller, make sure its array of inputs matches or exceeds the number of source components in your system. Your system is likely to expand in the future, so look for a controller with at least one more input than you need right now.

All controllers have inputs for digital audio signals as well as for analog. These inputs receive the digital-audio output of a Blu-ray Disc player, DVR, or music server. The signals carried on these digital connections include Dolby Digital, DTS, Dolby Surround, and 2-channel PCM (pulse code modulation) signals. These inputs are in addition to the HDMI inputs which have replaced a wide range of video and audio connections found on previous-generation home-theater equipment. HDMI carries high-definition video along with high-resolution multichannel audio on a single cable. Most controllers have at least four High-definition Digital Multimedia Interface (HDMI) inputs along with one or two HDMI outputs, the latter to feed your video display. Modern controllers can pass Ultra HD 4k signals in addition to conventional HD signals via HDMI. Note that a standard coaxial or TosLink optical connection cannot carry the newer surround-sound formats (Dolby Digital Plus, Dolby TrueHD, DTS HD, DTS HD Master Audio). These formats can be carried only on HDML In fact, component video, composite video, and S-video jacks are quickly becoming things of the past as they are supplanted by HDMI. Virtually all televisions, Blu-ray Disc players, A/V receivers and controllers, and set-top boxes now use the HDMI interface. (An HDMI cable is shown in [Fig.11-3.](#))



Fig.11-3 An HDMI connector (Courtesy Monster Cable Products)

Unlike A/V receivers, with their RCA jacks for inputs and outputs, higher-end controllers sometimes have balanced inputs and outputs. A balanced cable, described in [Chapter 11](#), is found only on higher-end controllers and power amplifiers. An RCA cable, also called an unbalanced or single-ended cable, carries the signal on one conductor plus a ground conductor. A balanced cable carries the signal on two conductors plus a ground. Balanced cables are used in professional audio, and have some advantages over unbalanced cables. Controllers with balanced inputs and outputs will usually offer both balanced and unbalanced connection jacks to ensure compatibility with a wide range of power amplifiers and source components. To use a controller's balanced inputs, your source components must also have balanced output jacks. And a controller's balanced output jacks will work only with a power amplifier with balanced input jacks. If your source components or power amplifier lack balanced jacks, you can still connect them through the unbalanced (RCA) connections. The controller shown in [Fig.11-2](#) has balanced outputs on all channels.

Automatic Calibration:

A controller must be calibrated for the particular loudspeakers and room with which it is used. Specifically, you must adjust the levels of each channel so that they are the same loudness, as well as tell the controller how far you are sitting from each speaker so that the controller can impose a delay on signals feeding the speakers closest to the listener. This last provision ensures that the sounds from all of the speakers reach your ears at the same time. Some controllers, however, provide automatic calibration that adjusts the level and delay of each channel with the push of one or two buttons. You position the supplied calibration microphone at the listening seat, press a button or two, and the controller does the rest. The controller emits a series of test signals that are reproduced by your speakers, picked up by the microphone, and analyzed by the controller. The controller is now precisely calibrated for your speakers, their positions, and your listening position.

DSP Speaker and Room Correction:

Many controllers have the ability to measure the frequency response of your loudspeakers and room, and then apply a correction in the digital domain to deliver perfectly flat frequency response at the listening position. As with automatic calibration, you position a microphone at the listening position and wait for the system to reproduce and analyze a series of test signals. If the system detects, for example, a peak in the frequency response of +4dB at 125Hz, the correction system creates a filter that attenuates the signal by 4dB at 125Hz. In essence, the DSP correction system “pre-distorts” the signal so that when it is reproduced by your speakers in your room, the result is flat frequency response. In practice, DSP correction systems are vastly more complex and sophisticated than this simple example suggests. The more sophisticated correction systems, such as Audyssey, can have a profound effect on the quality of the bass, making it smoother and less colored. DSP room-correction systems generally operate over the entire audioband, and can significantly change your system’s tonal balance, sometimes for the worse.

Analog Bypass Modes:

For the music lover shopping for a controller that will also serve as a 2-channel preamplifier for his system, one of the most significant considerations is its performance with 2-channel analog sources (especially if you have an extensive vinyl collection) and signals from a high-resolution disc player. In that case, you'll want a controller that has an analog bypass mode. Without an analog bypass mode, the analog signal will be converted to digital and back to analog as it passes through the controller. Digital conversion is far from sonically transparent, so the sound will suffer.

There are two catches to look out for regarding the bypass mode. First, the controller must have an analog volume control. Most modern controllers adjust the volume digitally in their DSP (digital signal processing) chips, which usually don't sound as good as an old-fashioned volume control. Second, whenever you engage bass management, even a controller with a bypass mode will convert the analog signal to digital because bass management is performed by the DSP chips. If you have a subwoofer with satellite speakers and use the controller's crossover to divide the frequency spectrum, the bypass mode won't remove the A/D and D/A conversions from the signal path. This is a serious limitation for music lovers who demand the ultimate in sound quality.

One novel approach used by a few manufacturers is to split the incoming analog signal and filter bass from the left and right outputs with an analog filter. This signal, which has never been digitized, drives the left and right channels. The other half of the split signal is digitized and processed through the controller's bass management, which derives a subwoofer signal. It's not pure analog bypass because the signal driving the subwoofer has been digitized, but bass frequencies are less prone to sonic degradation by digital conversions.

If you've found a controller that meets your home-theater needs but lacks an analog bypass mode, you can still enjoy uncompromised music performance by adding an analog preamplifier to your system. Your analog source components connect to the analog preamplifier, as do the left and right outputs from the controller. In essence, the preamplifier is inserted in the signal path between the controller and left- and right-channel power amplifiers. When you play movies through the controller, the analog preamp acts as if it's not there, and has no effect on the signal. Some preamps have a feature (called something like "theater pass-through") that sets their gain (volume) at a predetermined level when you play movies, so that you don't have to adjust the preamp's volume control to

achieve correct volume from the left and right speakers in relation to the other speakers. This setup is explained in detail later in this chapter, under “Adding Home Theater Without Compromising Music Performance.”

Bass-Management Flexibility:

An important controller function is bass management, the subsystem that lets you selectively direct bass information in the soundtrack to the main loudspeakers or to the subwoofer. Bass management allows a controller to work correctly with a wide variety of speaker systems. For example, if you have five small speakers and a subwoofer, you tell the controller to filter bass from each of the five channels and to direct it, in sum, to the subwoofer. With a Dolby Digital or DTS movie, the bass from the left, center, right, and surround channels is mixed with the LFE channel to drive the subwoofer. The controller's "subwoofer" output is the LFE channel selectively mixed with bass from any of the other channels. The bass management in most controllers lets you direct the full audioband to the left and right channels (including the LFE channel), but filter bass from the center and surround channels.

A feature of the most advanced controllers is their ability to specify the crossover frequency and slopes between the subwoofer and main speakers. The crossover is implemented in the digital domain with DSP. That is, the audioband is divided by performing mathematical computations on the ones and zeros representing the audio signal. Splitting the audioband into bass and treble in the controller is a vastly better approach than subjecting the analog audio signal to the capacitors, resistors, and inductors found in the crossovers built into subwoofers.

The most sophisticated controllers let you specify the crossover frequency (40Hz, 80Hz, 120Hz, for examples), as well as the crossover slope or phase characteristics. The greater the flexibility in specifying crossover frequency, slope, and phase, the greater the likelihood that you can seamlessly integrate the subwoofer's output with that of the main speakers. (Crossover slope and phase are described in [Chapter 6](#).)

Keeping low bass out of smaller loudspeakers confers large advantages in the speakers' power handling, dynamic range, midrange clarity, and sense of ease. When the woofer doesn't have to move back and forth a long distance trying to reproduce low bass, the midrange sounds cleaner and the speaker can reproduce loud peaks without distortion.

8-Channel Analog Input:

If you plan to use a multichannel SACD player with a controller, you'll need a feature called "discrete multichannel analog input." This is an input consisting of six or eight RCA jacks that accept the six outputs from a multichannel SACD player. This input bypasses all digital processing and sends the multichannel signal to the controller's multichannel outputs. Note that SACD players provide bass-management within the player, obviating the need to engage the controller's DSP-based bass management.

Digital-to-Analog Conversion:

Built into every 7.1-channel controller are eight digital-to-analog converters (DACs) and eight analog output stages. The DACs convert the digital data for each channel into analog signals. The quality of these DACs and the subsequent analog output stage (which drives the power amplifier through interconnects) is crucial to realizing good sound quality. DACs vary greatly in their sound, and a poor-sounding DAC (or a poor implementation of a good one) can ruin an otherwise excellent controller. More expensive controllers use higher-quality parts and design techniques, including metal-film resistors, polystyrene capacitors, four-layer circuit boards, and exotic circuit-board material. Also look for analog stages made from discrete transistors instead of inexpensive operational-amplifier (op-amp) chips. Virtually all of today's controllers are compatible with high-resolution signals, up to 192kHz sampling rate and 24-bit word length. Some high-end companies now have considerable expertise in designing cutting-edge digital converters, expertise they can apply to building multichannel digital controllers. Much of what you're paying for in a high-end controller is the quality of its DACs and analog output stages.

7.1- and 9.1-Channel Playback from 5.1-Channel Sources:

Some home-theater receivers and controllers provide 7.1- or 9.1-channel playback from 5.1-channel sources such as Dolby Digital and DTS. These products have eight or ten channels (seven or nine channels plus a subwoofer) rather than the six channels of most Dolby Digital- and DTS-equipped products. A 7.1-channel system can use the two additional channels (beyond the essential left, center, right, surround left, surround right) as “height” channels in a system with Dolby Pro Logic IIz decoding, or to drive two surround-back speakers placed behind the listener, augmenting the two surround speakers at the sides of the listening position. A 9.1-channel system can accommodate height speakers as well as surround-back speakers.

Making effective use of the additional surround speakers requires sophisticated signal processing in the controller. Controllers with THX Ultra2 processing incorporate a circuit that creates a 7.1-channel signal from any source, allowing you to use your entire 7.1-channel speaker system with all source material, not just those movies that have been mixed in 7.1-channel sound. Several Dolby and DTS processes (described at the end of this chapter) are also designed to convert 5.1-channel sources into 7.1- or 9.1-channel signals for reproduction by a 7.1- or 9.1-channel loudspeaker array.

THX Certification:

Some controllers are THX-certified, meaning that they incorporate certain types of signal processing to better translate film sound into a home environment and meet a set of specific technical criteria, established by Lucasfilm, do just that. THX processing in a controller includes re-equalization, surround decorrelation, and timbre matching. *Re-equalization* removes excessive brightness from film soundtracks for home-theater playback. *Surround decorrelation* improves the sense of spaciousness from the surround channels by introducing slight differences between the two surround signals. *Timbre matching* compensates for the fact that the ear doesn't hear the same tonal balance of sounds in all directions. The THX timbre-matching circuit ensures that sounds from the front and rear have the same timbre (tonal balance). To keep their products' prices low, some manufacturers of budget controllers license from Lucasfilm only the THX re-equalization circuit. Such controllers cannot, however, be billed as "THX-certified."

A THX-certified controller will also include a crossover circuit to separate the bass from the rest of the audioband when a subwoofer is connected to the system. The crossover keeps bass out of the front-channel signals, and filters mid- and high frequencies from the subwoofer output. A THX crossover has a specified cutoff frequency of 80Hz. (The cutoff frequency is the frequency below which bass is sent to the subwoofer and above which higher frequencies are sent to the front-channel speakers.)

3D Capability, Network Connection, Wireless Streaming, and Multi-Zone:

If you think you'll want 3D video in your home, choose a controller that can accommodate 3D signals, which requires HDMI 1.4a or later connectivity. In addition, the Audio Return Channel feature of HDMI 1.4a simplifies your system connection by allowing an HDMI cable to send the audio signal back from the video display to the controller. This feature obviates the need for connecting an audio cable from your video display back to your controller. If you're watching TV from a tuner source, for example, the audio signal from the TV will return to the controller via the HDMI cable. All HDMI cables support Audio Return Channel functionality.

Some controllers provide an Ethernet jack for networking with other devices in the home. This feature allows you to access content (photos and music files, for examples) stored on a PC or other device on the network, such as a network-attached storage (NAS) drive. The Ethernet port can be used for software updates, as well as for accessing streaming media such as Pandora and Rhapsody. Today's controllers also offer wireless connectivity to a network, along with the ability to be controlled by apps running on mobile devices. All the wireless and networking features described in [Chapter 8](#) can be found in modern controllers and audio/video receivers.

Most controllers have multi-zone capability, which allows you to watch one source in one room and to another source in a second room. More elaborate versions of this feature provide greater control over the Zone 2 signal. Keep in mind that unless the controller has wireless capability, you'll need to run wires to the second zone.

Controllers will sometimes offer 9.1-channel capability, with the additional two channels (beyond left, center, right, surround left, surround right, surround back left, surround back right) for reproducing the additional height or width channels offered by Audyssey DSX or the height channels in Dolby Pro Logic IIz.

A front-panel USB port allows you to access content on a USB stick, such as music files, photos, and album art.

Multichannel Power Amplifiers

The power amplifier is the workhorse of your home-theater system. It takes in line-level signals from the output of your controller and amplifies them into powerful signals that can drive your home-theater speaker system. A power amplifier has line-level inputs to receive signals from the controller, and terminals for connecting speaker cables. The power amplifier is the last component in the signal path before the speakers.

A separate multichannel power amplifier can offer higher output power than the power amplifiers found in A/V receivers. The most powerful receivers can have outputs of 140 watts per channel; separate power amplifiers offer as much as 350Wpc.

A power amplifier is described by the number of speakers it can amplify. The most common home-theater power amplifiers have five channels, for powering the left, center, right, surround left, and surround right speakers. The next most common number of channels is three. A 3-channel amplifier can power the front three speakers, leaving a 2-channel amp to handle the two surround speakers. If you already have a stereo (2-channel) amplifier, you can update your system for home theater by adding a 3-channel home-theater amplifier. Alternatively, a 3-channel amplifier can power the center and two surround speakers, leaving your stereo amplifier to drive the left and right speakers. As the Dolby Digital Plus and DTS HD formats become more popular, so do power amplifiers with seven channels. Note that if you upgrade your controller to a 7.1-channel model, you can augment your 5-channel power amplifier with a 2-channel amp to power the surround-back speakers, or the front height speakers (if your receiver or controller offers Pro Logic IIz decoding).

Some home-theater power amplifiers, called monoblocks, drive only one amplifier channel per housing. The monoblock approach lets you upgrade more easily by buying power for only as many channels as you need. Monoblocks have the potential of better sonic performance, but also cost more than the same number of amplifier channels in a multichannel power amplifier.

How to Choose a Home-Theater Power Amplifier

High-quality multichannel power amplifiers are big, heavy, and can be expensive. The more output power provided by the amplifier, the bigger, heavier, and more expensive it will be. Most home-theater enthusiasts listen at moderate levels to loudspeakers of average sensitivity (85–88dB) in rooms of average size (3000–4500 cubic feet) with average furnishings (carpet, some drapes, a stuffed sofa), and so will need a minimum of about 70W for each of the five, six, or seven channels.

Some multichannel power amplifiers feature a single large power transformer to supply all five, six, or seven channels. Others provide smaller, separate transformers, one for each channel. There is no consensus as to which approach is better. The argument for a single large transformer goes like this: For 2-channel music sources, the amplifier benefits from the extra power-supply capacity provided by a huge transformer. Moreover, even when multichannel sources are played, the rear channels don't work as hard as the front, allowing some of their power-supply capacity to be delivered where it's needed. The opposing point of view holds that separate, smaller transformers completely isolate the channels from each other and eliminate interaction between channels. For example, a loud explosion in the film soundtrack reproduced by the center channel can pull power from the transformer needed for the musical score being reproduced by the left and right channels.

Note that driving five, six, or even seven speakers places extraordinary current demands on an amplifier. There's a way to reduce the strain placed on multichannel amplifiers: use high-sensitivity speakers. Remember that every 3dB increase in sensitivity is equivalent to doubling amplifier power.

More specific information on power amplifiers can be found in [Chapter 5](#). The discussion of sound quality, build quality, and matching to speaker sensitivity applies as much to multichannel amplifiers as to 2-channel amplifiers.

Loudspeakers for Home Theater

A home-theater loudspeaker system will perform double duty, reproducing 2-channel music sources through the front left and right speakers, and film soundtracks through all five, six, or seven speakers (plus an optional subwoofer). Home-theater speakers can be evaluated using the criteria described in [Chapters 3 and 6](#), with one difference: film soundtracks present somewhat different challenges to a loudspeaker system than does music. Specifically, film soundtracks are generally more dynamic than most music, and have much more low-frequency energy. Consequently, great home-theater sound requires speakers that not only meet our criteria for good music reproduction, but also have a wide dynamic range and can handle loud, complex signals without sounding smeared or confused.

The Center-Channel Speaker

Though the center-channel loudspeaker carries a large part of the film soundtrack—nearly all the dialogue, many effects, and some of the music—only recently have center-channel speakers been elevated from afterthoughts to recognition that the center speaker is the anchor of the entire home-theater loudspeaker system.

The center speaker is usually positioned horizontally above or below the video monitor, either attached to the wall, mounted in the wall, or, if you're using a front-projection system, behind an "acoustically transparent" screen.

Because a stereo system uses only two speakers—left and right—across the front, you may be wondering why you need a third speaker between them for home theater. The center-channel speaker provides many advantages in a home-theater system. First, it anchors dialogue and other sounds directly associated with action on the screen in the center of the soundstage. When we see characters speaking, we want the sound to appear to come from their visual images. Similarly, when sounds are panned (moved from one location to another) across the front, we want the sound to move seamlessly from one side to another. For example, if the image of a car travels from the left side of the screen to the right, the sounds of the car's engine and tires should travel with it, precisely tracking the car's movement. Without a center speaker, we may hear a gap in the middle as the car sounds jump from the left speaker to the right. The center speaker makes sure on-screen sounds come from the screen.

A 2-channel stereo system is capable of producing a sonic image between the two speakers. This so-called "phantom" center image is created by the brain when the same signal is present in both ears. A sound source in front of us in real life produces soundwaves that strike both ears simultaneously. The brain interprets these cues to determine that the sound source is directly in front of us. Similarly, two speakers reproducing the same sound send the same signal to both ears, fooling the brain into thinking the sound is in front of us.

For two speakers to create this phantom center image, they must be precisely set up, and the listener must sit exactly the same distance from each of them—if you sit off to the right, the center image will pull to the right. In addition, creating a phantom center image from two loudspeakers may be a learned skill; some listeners may never hear a phantom center image. There's also evidence that the work the brain must perform to conjure up this phantom center image is distracting and fatiguing.

These problems are overcome by putting a center-channel speaker between the left and right speakers. Dialogue and onscreen sounds are firmly anchored on the screen for all listeners, not just those sitting in the middle. With three speakers across the front, someone sitting way off at the left end of the couch can still hear dialogue coming from the area of the screen—not just from the left speaker. The center speaker also prevents the entire front soundfield from collapsing into the speaker closest to where you're sitting. Moreover, the center speaker provides a tangible sound source directly in front of you; your brain doesn't have to work to create a phantom image between the left and right speakers. Finally, the center speaker reduces the burden on the left and right speakers. With three speakers reproducing sound, each can be driven at a lower level, for cleaner sound.

Adding a Center Speaker to Your System

Although you can add a different brand of center speaker to your existing left and right speakers, you're better off with three matched speakers across the front of your home theater. Speakers all sound different from one another. No matter how good the quality of a speaker, it will have some coloration, or variations from accuracy, in its sound. I mentioned earlier that you want to hear a seamless movement of sounds across the front soundfield. If the center speaker has a different sound from the left and right speakers, you'll never achieve a smooth and continuous soundfield across the front—when sound sources move from one side to another through the center speaker, the sound's character will abruptly change.

The best solution is to buy three matched front-channel speakers. Their identical tonal characteristics will not only provide smooth panning of sounds, but also produce a more stable and coherent soundfield across the front of the room. If you already have high-quality left and right speakers that you want to keep, buy a center-channel speaker made by the manufacturer of the left and right speakers. They probably won't be as well matched as a three-piece system, but the added center speaker is much more likely to sound similar to your existing left and right speakers. Separate center speakers sometimes use the same drivers (the raw speaker cones themselves) and other parts as the stereo speakers from the same manufacturer.

For state-of-the-art home theater, the three front speakers should be identical. Although most home-theater enthusiasts will use a smaller center speaker that will fit above or below a flat-panel television, and doesn't reproduce much bass, more ambitious systems use large, full-range center speakers. These speakers are hard to position, and are generally used only with front-projector systems; the center speaker can be placed behind a perforated projection screen, just as in a movie theater.

Left and Right Speakers

The left and right loudspeakers reproduce most of the film's musical score and many of the effects. And if you aren't using a subwoofer, nearly all the bass will be reproduced by the left and right speakers. Consequently, left and right speakers are often the largest speakers in a home-theater system.

Left and right speakers can also be small units that sit on speaker stands or in an entertainment cabinet. If that's the case, the small speakers must be used with a subwoofer to reproduce bass. The small left and right speakers reproduce the midrange and treble frequencies, and the subwoofer handles all the bass. Such a system, called a satellite /subwoofer system, is ideal if your available space is limited, or if you want the left and right speakers to better blend into your decor. Satellite speakers can be small and unobtrusive, and the subwoofer can be tucked out of the way. If you build the three front speakers into a wall in a custom installation, they'll most likely be satellites.

Surround Speakers

The surround loudspeakers are completely different in design and function from left, center, and right speakers. Their job is to re-create a diffuse atmosphere of sound effects to envelop us in a subtle sonic environment that puts us in the action happening onscreen. Unlike front and center speakers which anchor the sound onscreen, surround speakers should “disappear” into a diffuse wash of sound all around us.

A good example of how surround speakers create an atmosphere comes from the film *Round Midnight*. Toward the beginning of the movie, the character François is outside a Paris jazz club in a driving rainstorm, listening to Dexter Gordon’s character playing inside. The scene cuts between the intimate sound of the club and the rainy Paris street. Inside the club, the sound is direct and immediate, to reflect the camera’s perspective of just a few feet from the musical group. When the scene cuts to the street, we are surrounded by the expansive sounds of rain, cars driving by, people talking as they walk past—in other words, all the ambiance and atmosphere of a rainy night in Paris. This envelopment is largely created by the surround speakers. We don’t want to hear the rain and street sounds coming from two locations behind us, but to be surrounded by the sounds, as we’d hear them in real life. The surround speakers perform this subtle yet vital role in home theater.

Surround speakers envelop us by their design and placement in the home-theater room. They are best located to the side or rear of the listening position, and several feet above ear level. Because they don’t have to reproduce bass, surround speakers can be small and unobtrusive, and are often mounted on or inside a wall.

Dipolar and Bipolar Surround Speakers

Most of the surround speakers' ability to wrap us in sound comes from their dipolar design. *Dipolar* simply means that they produce sound to the front and rear equally. While front speakers have one set of drivers that project sound forward, dipolar surround speakers have two sets of drivers, mounted front and back. This arrangement produces a directional pattern that fires to the front and back of the room (see [Chapter 6](#)). Because the surround speakers are positioned to the sides and fire to the front and back of the room, we hear no direct sound from them. The listener sits in the surround speakers' null—the point where they don't directly project sound. Instead of hearing direct sound from the surround speakers, we hear their sound after it's been reflected by, or bounced off, the room's walls and furnishings. The surround speakers' dipolar directional pattern makes their sound diffuse and harder to localize. You shouldn't be able to tell where a properly set-up surround speaker is just by listening.

You may see surround speakers called either dipolar or *bipolar*. Although both produce sound equally from the front and rear, the bipolar speaker produces sound from the rear that is in phase with the front sound. A dipolar produces sound from the rear that is out of phase with the front sound. In phase means that the front and rear waves have the same polarity: when the front-firing woofer moves forward to create a positive pressure wave, the rear-firing woofer also moves forward. The opposite is true in a dipole: when the front-firing woofer moves forward, the rear-firing woofer moves backward. In a dipole, the front and back waves are identical, but have the opposite polarity. Both bipolar and dipolar achieve the objective of enveloping the viewer-listener in the film's aural ambiance, but the dipolar type is preferred because it creates a greater null, or reduction in sound, at the side of the speaker that faces the listener.

A true dipolar or bipolar surround speaker can't be flush-mounted inside a wall. Instead, it must be mounted *on* the wall so it can radiate sound toward the front and rear of the room. Some surround speakers approximate a dipolar radiation pattern by having small drivers mounted on an angled baffle. Although not true dipoles, these flush-mounted surround speakers accomplish the goal of creating a diffuse ambiance and fit unobtrusively in your living room.

The amount of sound produced by the surround speakers is adjustable through the receiver's or controller's remote control. Their volume level should be set so that you never hear them directly; but if you turned them off, the soundfield would collapse toward the front. All too often, consumers set

surround speaker levels too high; their thinking seems to be that if they bought the speakers, they should be able to hear them.

Surround-Back Speakers

As described earlier, the 7.1-channel formats add two additional surround speakers behind the listener. According to Lucasfilm, the four surround channels are best reproduced by four identical dipolar speakers. The two surround speakers remain at their normal positions (normal for 5.1-channel reproduction), which is 110° , or slightly behind the listener on the side walls. The surround back speakers are positioned at 150° , as shown in Fig.12-4 later in this chapter.

The two surround-back speakers should be arranged so that their facing sides are in phase with each other. That is, the side of the dipole that would face the back of the room if mounted on a sidewall should face toward the room's centerline when the speaker is mounted on the rear wall. This arrangement keeps the outputs of the drivers of two surround-back speakers that face each other in phase, contributing to their ability to create an image between them directly behind the listening position.

Height Speakers

The height channels in Dolby Pro Logic IIz attempt to create a more immersive experience by reproducing sound from above the plane of the front loudspeakers. They should be positioned at least three feet directly above the front left and right speakers. The height speakers reproduce ambient sounds such as rain rather than music and effects. Consequently, the height speakers can be smaller than the left and right speakers.

Subwoofers

A subwoofer is a speaker that reproduces only low bass. Subwoofers are usually squarish or rectangular cabinets that can be positioned nearly anywhere in the home-theater room. A good subwoofer adds quite a bit of impact and visceral thrill to the home-theater experience.

Most subwoofers for home-theater use are powered or active subwoofers. These terms describe a subwoofer with a built-in power amplifier that drives its large woofer cone. The powered subwoofer takes a line-level output from the jack marked Subwoofer Out on your A/V receiver or controller, and converts that electrical signal into sound. Because powered subwoofers have built-in amplifiers, they must be plugged into an AC wall outlet. (See [Chapter 6](#) for more on subwoofers.)

The subwoofer is typically driven by the LFE Output jack on an AVR or controller. The LFE output is a monophonic mix of the LFE channel along with the bass from any channel designated as “Small” during the controller’s initial setup.

Setting Up a Home Theater

How a home-theater system is set up and calibrated has an enormous influence on the quality of the sound you hear. Just as with a 2-channel music system, ideal loudspeaker placement and careful attention to detail will reward you with superior sound. The acoustic principles described in [Chapter 14](#) apply equally to a multichannel audio system.

Basic Setup

The first step in setting up your system is deciding where to install it. More than likely, the room where you watch television will become the home-theater room. But this room may have been arranged for casual TV watching, not for experiencing the full impact of what a home-theater system can deliver. Now is the time to rethink the room's layout and consider making the home-theater system the room's central focus.

The home-theater room should be arranged with the video monitor and front loudspeakers at one end, the couch or viewing chairs at the other. Ideally, the couch will be positioned partway into the room and not against the back wall. This location is better acoustically (excessive bass builds up close to the back wall), and allows the surround speakers to better do their job of enveloping you in the film's soundtrack. Of course, the setup should provide room for the left and right speakers on either side of the television.

The optimal viewing distance between you and the video display is determined by the video display's size. The smaller the display, the closer you should sit to maintain the maximum image area in your field of view. A good rule of thumb is to sit three to four picture heights away from the screen. You should also try to sit no more than 15° off to the side of the video monitor. Plasma panels and LCD displays, with their wide fields of view, maintain a good picture even when watched from far off to the side.

Acoustical Treatment

Before we get to the specifics of loudspeaker placement, here are a few guidelines on how to make your home-theater room sound better. First, if you have a bare tile or wood floor between you and the three front speakers, covering the floor with a rug is probably the single most important thing you can do to improve your room's acoustics. The rug absorbs unwanted acoustic reflections from the floor, which improves dialogue intelligibility and adds to a sense of clarity. Similarly, absorbing the acoustic reflections from the sidewalls between you and the left and right speakers also helps you get better sound. Bookcases and hanging rugs are both effective in absorbing or diffusing sidewall reflections. Generally, the front of the room should be acoustically absorptive, with drapes, carpets, and hanging rugs.

Conversely, the back of the room should be more acoustically reflective. Bare walls behind the listening/viewing position tend to reflect sound and add to the feeling of spaciousness generated by the surround speakers. The sounds produced by a movie theater's array of surround speakers all reach your ears at different times, increasing the feeling of envelopment. But because a home-theater system uses just two surround speakers, we must help them achieve their goal by providing reflective surfaces near them.

Ideally, the lower portion of the room's front end should be acoustically absorptive, the upper portion acoustically reflective. This situation absorbs unwanted wall reflections from the three front speakers, but lets the dipole surround speakers positioned on the sidewall just below the ceiling create a bigger sense of envelopment by reflecting sound off the front walls as well as the rear walls.

Speaker Placement

Where speakers are placed in your room has a huge influence on the quality of sound you will achieve. Correct speaker placement can make the difference between good and spectacular performance from the same equipment. If you follow a few simple speaker-placement techniques, you'll be well on your way toward getting the best sound from your system. Correct speaker placement is the single most important factor *in achieving* good sound.

Surround Speaker Placement

Let's start with the most difficult of a home-theater system's six speakers: the two surround speakers. As I said earlier, the surround speakers are usually dipolar. That is, they produce sound equally to the front and the rear of the speaker cabinet. Dipoles produce a null, or area of no sound, at the cabinet sides. Here's a simple rule for positioning dipolar surround speakers, no matter the installation: point the surround speakers' side panels, not the drivers, at the listening position. Never point the surround speakers directly at the listening position.

Ideally, the surround speakers would be placed against the side walls directly to the sides of the listening position, and at least two feet above the listeners' heads when seated. Dipolar surround speakers should never be positioned in front of the listening seat, where they could be heard directly; the surround speakers' sound should be reflected around the room before reaching the listener. This helps to mimic the sound of a movie theater's surround array with just two speakers, and to create the envelopment that's so important to the home-theater experience.

Keep in mind that reflections of the sounds from the front speakers are bad; reflections from the rear speakers are good. Well-positioned surround speakers make a tremendous difference in the quality of the home-theater experience by re-creating the sense of ambience and space we hear in a movie theater.

The surround-back speakers should be positioned on the rear wall, behind the listeners, at the positions shown in [Fig.11-4](#). This placement provides the greatest sense of envelopment, along with the most continuous soundfield behind the listener. The two surround-back speakers should be positioned so that the speakers facing the room's centerline are in-phase with each other. This produces a central image between the two surround-back speakers.

Center Speaker Placement

Correct positioning of the center speaker results in better dialogue intelligibility, smoother movement of onscreen sounds, and a more spacious soundstage. As just mentioned, reflections from the front speakers that reach the listener work against good sound. This is particularly true of the center speaker, which carries the all-important dialogue. Because most center-channel speakers are mounted above or below a flat-panel television—which is highly reflective—there’s the potential for unwanted reflections.




Fig.11-4 The surround loudspeakers should be positioned at 110° to the listening position and the surround-back speakers at 150°.

To reduce reflections of the sounds from the center speaker, align its front edge absolutely flush with the front of the television, so that they form as contiguous a surface as possible. This positioning reduces a phenomenon called *diffraction*, which is the re-radiation of sound when the soundwaves encounter a discontinuity. (Specifically, diffraction introduces variations in the speaker's frequency response.) Think of waves rippling in a pond: If you put a stick in the

water, the waves will strike the stick and be re-radiated around it. The same thing happens when the sound from the center speaker encounters the television: the television re-radiates that sound toward the listener, making the overall sound less accurate. That's why making the center speaker flush with the front of the television is so important: such positioning gives the soundwaves only a single effective surface to radiate from.

High-end loudspeaker manufacturers carefully design their enclosures to reduce diffraction. You'll notice that high-quality loudspeakers have smooth surfaces around the speaker cones. Some have rounded fronts, and even the mounting bolts securing the individual drivers to the cabinet are recessed, to reduce diffraction.

You should also tilt the center speaker down toward the listening position if the center speaker is mounted above the television. If the center speaker is located below the television, point it up toward the listeners. Direct the center speaker's sound toward the height of the listeners' ears when seated. Some manufacturers of speaker stands offer short stands that are angled, so that the center speaker is tilted up toward the listening position. Finally, be sure the difference in height between the center and left and right speakers is less than 2', measured from each speaker's tweeters.

Left and Right Speaker Placement

The left and right speakers should be pulled slightly in front of the television, and not be in a straight line with it (as some owner's manuals erroneously advise). Positioning the left and right speakers so that they form a gentle arc with the center speaker has two advantages. First, acoustic reflections from the television are reduced. Second, the left, right, and center speakers should now be all the same distance from the listener, so that the sounds from each of the three front speakers will reach the listener at the same time. If the three front speakers were all on the same plane, we would hear the center speaker slightly ahead of the sounds of the left and right speakers, which tends to make sounds bunch together in the center and reduces the soundstage's width.

If you're using large, full-range left and right speakers, putting the speakers near the front and sidewalls will increase the amount of bass you hear. The walls reinforce bass to produce a weightier sound. This can be a good or a bad thing, depending on the speakers and your room. Speakers that already produce lots of bass will probably sound boomy and unnatural if put too close to the front wall. Pulling them out into the room often produces huge improvements in bass definition (the bass no longer sounds like a continuous blur) and midrange clarity. The unnatural boominess of a speaker placed too close to a wall makes the midrange sound thick, and lacking in openness and clarity. If you're using small left and right speakers with a separate subwoofer, the left and right speakers' placements in relation to the front wall are less important.

How far apart the left and right speakers are from each other is also important to achieving a natural and spacious sound. If the left and right speakers are too close to the video monitor, you'll never get a sense of large space at the front of the room. If they're too far apart, onscreen sounds are reproduced too far away from their apparent sources on the screen. It can be subconsciously confusing and distracting to hear a sound come from a place different from the visual position of its source. A good rule of thumb is to position the left and right speakers so that each forms a 60° angle with the listening/viewing position.

Left and right speakers can be positioned so that they point straight ahead, or are angled in (called toe-in) toward the listening position. No toe-in (speakers pointing straight ahead) produces a wider and more spacious soundstage, but at the expense of precise localization of sounds. Too much toe-in creates a more tightly focused sound but a narrower soundstage. Toe-in has the advantage of directing more sound at the listener and less sound toward the sidewalls.

Reflection of sound from the sidewalls is the single worst problem in getting good sound from your speakers.

Some speakers are meant to be toed in; these positions provide their flattest (most accurate) frequency response. Others achieve their best performance with no toe-in. If you toe in speakers designed for no toe-in, the sound will be too bright. Consult your owner's manual or your dealer for guidance. You can also experiment by playing the same section of a movie soundtrack or a piece of music with varying degrees of speaker toe-in to determine which position sounds best.

If you're using height speakers, position them directly above the left and right speakers. They should be at least 3' above those speakers.

Calibrating a Home Theater

Once your system is connected, you'll need to configure it for your particular room and speaker system. This is done through a series of menus or graphical icons generated by the A/V receiver or A/V controller and appearing on your video monitor. You use the receiver's remote control to make selections from the onscreen display. Every A/V receiver or controller includes instructions for setting up your system; the following section explains why you make certain selections, and in what circumstances.

When you turn on your system for the first time, don't be alarmed if you don't hear any sound when you expect to. (This usually happens when I first set up a product for review.) Don't turn up the volume to a high level and then start pushing buttons—when you push the right button, the sound will blast you out, possibly damaging your speakers, your amplifier or receiver, and/or your hearing. Turn the volume to a low level and make sure the correct source is selected. Check to make sure the Mute and Tape Monitor buttons aren't pressed. The lack of sound is probably due to one of these three incorrect settings.

Bass Management

First, you must tell the A/V receiver or controller what kind of loudspeakers you have, so it can direct bass appropriately. Most of today's controllers ask you whether each of the speakers is Large or Small. By selecting Small for the left and right speakers, you're telling the receiver to keep low-bass signals out of those speakers. This configuration is used when you have a subwoofer connected to the system to reproduce low bass instead of the left and right speakers. You'll also need to answer Yes when asked if the system includes a subwoofer. The bass-management option labeled "THX" in THX-certified products automatically sets all speakers to Small and engages the subwoofer.

If you use full-range left and right speakers and no subwoofer, answer Large in the setup menu when asked if the left and right speakers are Large or Small. The Large setting directs bass to the left and right speakers. The setup menu will also ask if the surround speakers are Large or Small. Nearly every installation will use the Small setting. Only if you have full-range, floorstanding speakers should you answer Large.

Setting Individual Channel Levels

Next you'll need to individually set the loudness of each of the six or eight channels (assuming you're using a subwoofer). In addition to providing an overall volume control, all A/V receivers and controllers provide adjustment of the individual channel volumes. This process begins by turning on the "test signal," a noise-like sound generated by the A/V receiver or controller. The noise is produced by each speaker in turn. Ideally, the noise signal's volume should be the same for each speaker when you're sitting at the listening/viewing position. If it isn't, you can adjust the volume of each speaker independently using the receiver's remote control. Individually adjusting the channels lets you compensate for different loudspeaker sensitivities, listening-room acoustics, and loudspeaker placements.

Although setting the individual channel levels by ear will get your system in the ballpark, a more precise calibration can be achieved by using a sound pressure level (SPL) meter. Available from RadioShack for about \$30, an SPL meter lets you accurately calibrate your home-theater system. RadioShack offers digital and analog SPL meters; buy the easier-to-read analog type. Switch the meter to the c-weighted position, slow response, and set the knob to 70. Turn on your A/V receiver's test signal and set each channel's volume until the meter's display reads "5." This indicates that the test signal is being reproduced at a level of 75dB. Repeat this procedure for each channel.

A powered subwoofer usually has a knob on its rear panel for setting volume level. When you set the subwoofer's level, the meter will be hard to read because the indicator will be jumping around (it's reading a low frequency). Watch the meter for a few minutes as it jumps, to get an idea of what the average level is.

These settings will get you very close to the optimal volume for each channel, but you should use your ear and some well-recorded film soundtracks to make the final adjustments. If you find that dialogue is hard to hear, a 2dB boost in the center-channel level will help bring it out (although the soundstage will then be focused more in the center than spread out across the front of the room). If the bass is thumpy and boomy, turn down the subwoofer. Don't be afraid to adjust the volume of the subwoofer, center channel, or surround speakers—each film soundtrack is mixed differently. You may find yourself slightly adjusting the surround- and center-channel levels at the beginning of each movie.

Although your ears should be the final judge when setting the individual channel levels, calibrating your system first with the SPL meter (or by ear with the

test signal) will at least get you started from the right place. The most common mistakes in setting channel levels are a subwoofer level set too high, and too much volume from the surround speakers. The bass shouldn't dominate the overall sound, but instead serve as the foundation for music and effects. Many listeners think the more bass you hear, the more "impressive" the sound. In reality, a constantly droning bass is fatiguing, and robs the soundtrack of impact and surprise when the filmmakers *want* you to hear bass. Low frequencies are used as punctuation in a film soundtrack; by keeping the subwoofer level appropriate, you'll achieve a more accurate and satisfying sound than if the subwoofer is constantly droning away. Though it's understandable that you paid for this big box in your living room and you want to hear what it does, try to avoid the temptation to set the subwoofer level too high.

Similarly, listeners who have never had surround speakers in their homes think they should hear the surround speakers at all times. In truth, film soundtracks don't always contain signals in the surround channels; long stretches of a movie may contain nothing at all in the surrounds. But even when a signal is driving the surround speakers, you should barely notice their presence. Remember, the surround channels are there to provide subtle ambience and envelopment. If you're consciously aware of the surround speakers because they're set too loud, the illusion they're supposed to be creating is diminished. Just as it's a temptation to set a subwoofer's level too high, don't turn up the surround speakers to the point where you hear them. Note that you won't hear any sound from the surround speakers unless the source program has been Dolby- or DTS-encoded.

Adding Home Theater without Compromising Music Performance

As Editor-in-Chief of *The Absolute Sound*, my job involves evaluating cutting-edge, state-of-the-art 2-channel audio products. But my dedicated listening room for 13 years also served as a home theater and multichannel music system. Clearly, the performance of my 2-channel system could not be compromised by the presence of home-theater products. Here's what I did to integrate home theater into my music-playback system, along with some tips for adding home theater to your music system. (I now have separate stereo and theater systems.)

First, my video display is a front-projection system with a retractable, motorized screen. The projector is at the back of the room, and when I listen to music, the screen is rolled up into a small enclosure. With the screen lowered, soundstage depth is compromised, as is the precision of image placement. If you use a front projector, a motorized screen's ability to retract for music listening is a big benefit. If you must use a fixed screen, drapes that can be drawn across the screen when you play music are effective at preventing the screen from reflecting sound.

Those who use a flat-panel television are faced with the challenge of having a large, acoustically reflective object near the loudspeakers. As described in [Chapter 14](#), absorbing acoustic reflections at the loudspeaker end of the room is of paramount importance. To minimize the television's degradation of the sound, move the television back as far as possible, and bring the left and right loudspeakers forward. You can get an idea of how much the television is degrading stereo imaging by throwing a blanket over the TV to absorb its reflections. If you hear a big improvement with the blanket, consider keeping the blanket on hand for critical musical listening.

As mentioned earlier, an analog-bypass mode on a controller is absolutely essential if you want to listen to analog sources with the highest possible fidelity. This feature passes analog input signals through the controller without converting the signal to digital and then back to analog. Keep in mind, however, that even the very best multichannel controllers fall short of the performance standards set by high-end 2-channel preamplifiers, even in analogbypass mode.

If you want uncompromised musical performance, you'll need a controller and a separate preamplifier. The preamplifier should have a "theater pass-

through” mode that sets one of the inputs at unity gain (the output signal’s amplitude is the same as the input signal). The left and right outputs from your controller drive this unity-gain input, and the preamplifier is connected to the power amplifier in the usual way. When you watch movies, it’s as though your preamp isn’t there. When you listen to music, it’s as though your controller isn’t there. With this technique, shown in [Fig.11-5](#), your analog sources never go through the controller. If you choose this option, you don’t need a controller with analog bypass, except on the multichannel input.

If you’re adding home theater to your music system and are happy with your power amplifier, you should choose a 3-channel amplifier to drive the center and two surround speakers, or a 5-channel amplifier if you want a 7.1-channel system employing two surround-back speakers. If you’re starting from scratch, or have been thinking about upgrading your amplifier, consider a multichannel amplifier; some of today’s multichannel amplifiers are fully the equal of high-end 2-channel models.




Fig.11-5 A preamplifier with a “theater pass-through” mode avoids passing analog music signals through a digital controller.

Some subwoofers have separate processor and line inputs. The processor input (a single jack) accepts the LFE output from a controller; the line input is stereo, and is connected to a second output from your preamplifier. When watching movies, set the subwoofer’s input switch to Processor. When listening to music, set the sub’s input switch to Line. This assumes, of course, that you prefer listening to music with the subwoofer engaged. Unless you have small loudspeakers, you’ll probably leave the subwoofer off for music listening by leaving the switch set to Processor.

Addendum: Surround-Sound Formats Explained

In this section I provide a fuller explanation of the various surround-sound formats, for those readers interested in more technical detail. Even if you're not technically inclined, this section includes useful information about choosing among the vast array of decoding formats in today's AVRs and controllers.

We've already covered Dolby Pro Logic, Dolby Digital, and DTS in this chapter. Now let's look at variations on those formats.

In mid-2001, Dolby Laboratories made available a more sophisticated version of Pro Logic decoding, called Pro Logic II. The idea behind Pro Logic II was to create from 2-channel sources a listening experience similar to that of a discrete 5.1-channel digital format. With more consumers having 5.1-channel playback available to them for reproducing music sources, Pro Logic II attempts to deliver multichannel sound from 2-channel recordings, even those that have not been surround-encoded.

Found in most A/V receivers and controllers made after early 2002, Pro Logic II offers improved performance over its predecessor in several areas. First, Pro Logic II delivers full-bandwidth stereo surround channels, rather than the single, bandwidth-limited monaural surround channel of conventional Pro Logic decoding. This attribute provides a more enveloping soundfield, greater precision in the placement and pans (movements) of sounds behind the listener, and more natural timbres of sounds reproduced by the surround channels. In this respect, Pro Logic II approximates the experience of listening to a 5.1-channel discrete digital source.

Pro Logic II also uses more sophisticated "steering" circuits that monitor the level in each channel, and selectively apply attenuation (reduction in level) to prevent sounds in one channel from leaking into another channel. Pro Logic IIx, announced in late 2003, creates 7.1-channel playback from 2-channel and 5.1-channel sources. In 2009, Dolby introduced Pro Logic IIz, which calls for two "height" speakers to be placed in the front of the room directly above the left and right speakers. Pro Logic IIz extracts this height information from 5.1- or 7.1-channel sources and directs it to the height speakers, which can be small and unobtrusive. The idea is to create a more immersive soundfield by expanding the soundfield vertically.

Not to be outdone, DTS has developed a suite of surround-decoding formats that either enhance the DTS experience or provide decoding of non-DTS sources such as conventional stereo.

DTS Neo:6 Music and Neo:6 Cinema are decoding algorithms that convert stereo- or Dolby-Surround-encoded 2-channel sources into multichannel surround sound. Neo:6 Music leaves the front left and right channel signals unprocessed for the purest reproduction, and extracts center- and surround-channel signals from the 2-channel source. DTS recommends this mode for all 2-channel sources, such as CD and FM broadcasts.

Neo:6 Cinema is similar to Dolby Pro Logic II decoding, and can be used with Dolby Surround-encoded sources. Neo:6 Cinema has a much larger effect on the signal, rearranging the signal distribution among the front three channels. Both Neo:6 Music and Neo:6 Cinema synthesize a 7.1-channel signal from 2-channel sources.

These decoding algorithms are very useful to those, like me, who greatly enjoy musical performances on DVD and Blu-ray Disc. When a DVD gives me the choice of listening to the Dolby Digital or 2-channel Surround-encoded linear pulse-code modulation (LPCM) mix, I always opt for the 2-channel mix. Although I lose all the advantages of Dolby Digital (complete channel separation, for example), I hear smoother treble, less grainy instrumental and vocal textures, and a greater sense of space. That's because Dolby Digital uses 384,000 or 448,000 bits per second to encode all 5.1 channels; the linear PCM track uses 1.536 *million* bits per second to encode just two channels. When using Pro Logic II and DTS Neo:6 to decode these two channels, however, I still hear excellent spatial resolution, but with the more natural timbres made possible by the linear PCM track's much higher bit rate. Dolby TrueHD and DTS HD Master Audio, explained in a minute, don't force such choices; they deliver high-resolution multichannel audio without compromise.

In 2004, Dolby Labs and Lucasfilm jointly developed the Dolby Digital EX format. Dolby Digital EX encodes a third surround channel in the existing left and right surround channels in the Dolby Digital signal. This additional surround channel, called surround back, is decoded on playback, and the signal drives one or two speakers positioned directly behind the listener. Dolby Digital EX allows filmmakers to more precisely position and pan (move) sounds around the room. For example, the sound of an object moving directly overhead, from the front of the room to the back, tends to become smeared; the sound starts in the front speakers, then splits along the sidewalls as it moves toward the rear of the room. By adding a third surround channel directly behind the listener, such "flyovers"

can be made to sound more realistic. Note that an EX-encoded soundtrack is compatible with all Dolby Digital playback equipment because the addition of the surround-back channel is encoded into a conventional 5.1-channel Dolby Digital signal. Dolby Digital EX is not a 6.1-channel format; it is still 5.1. -channels, but with the additional channel encoded within the 5.1-channel datastream.

DTS's equivalent format is called DTS-ES. You'll also see it called DTS-ES Matrix, because the additional surround channel is matrix-encoded into the existing left and right surround channels. DTS also offers DTS-ES Discrete, in which the surround-back channel is a discrete channel in the ES-Discrete bitstream. DTS-ES Discrete is a true 6.1-channel format. The advantage of a discrete format is complete separation between channels. Matrix surround systems have limited channel separation.

Note that Dolby Digital and DTS are "lossy" formats, meaning that some information is lost in the encoding and decoding process. This loss of information is intentional so that the number of bits required to represent the signal can be dramatically reduced. Consequently, Dolby Digital and DTS have reduced fidelity to the source.

The massive storage capacity of the Blu-ray Disc format has allowed Dolby Laboratories and DTS to develop better-sounding surround-sound formats with higher bit rates than Dolby Digital and DTS. From Dolby Laboratories we have two formats, Dolby Digital Plus and Dolby TrueHD. Dolby Digital Plus is an extension of the conventional Dolby Digital we've been listening to for years on DVD and HD television shows, but Dolby Digital Plus offers more channels, and a higher bit rate for better sound quality. Where Dolby Digital was limited to a maximum bit rate of 640kbps (kilobits per second), Dolby Digital Plus allows scalable bit rates up to 6Mbps (megabits per second) on BD. In addition to more bits per second, Dolby improved the encoding algorithms in Dolby Digital Plus for better sound.

For the ultimate in sound quality, Dolby created Dolby TrueHD, which delivers high-resolution multichannel audio with perfect bit-for-bit accuracy to the source. With TrueHD, you will hear in your home sound quality identical to what the engineers heard in the studio, all else being equal. Specifically, TrueHD is capable of eight channels of uncompressed audio at 192kHz/24-bit. This is made possible by TrueHD's maximum bit rate (on BD) of 27Mbps. TrueHD decoding is an option, rather than a requirement, in Blu-ray players. The availability of a consumer release format that delivers uncompressed high-resolution multichannel audio is nothing short of a revolution in home

entertainment. Concert videos and musical performances on BD with uncompressed high-resolution multichannel audio are spectacularly great.

With more than 45 million Dolby Digital decoders in use throughout the world, Dolby made sure the new formats were backward-compatible with your receiver or controller. BD players output a conventional Dolby Digital or DTS bitstream on the familiar coaxial or optical digital outputs for connection to your receiver. This output will likely sound a little better than Dolby Digital from DVD because it always operates at the highest possible data rate of 640kbps (Dolby Digital on DVD is limited by the DVD format to 448kbps, but is typically 384kbps).

DTS has introduced its own high-resolution surround-sound audio format for BD, called DTS-HD. This new format is an extension of conventional DTS, offering scalable bit rates from conventional DTS at 754kbps, all the way to 6Mbps on BD. A variant of DTS-HD, called DTS-HD Master Audio, can operate at higher bit rates (up to 24Mbps on BD) for true lossless encoding of high-resolution multichannel audio. The “Master Audio” tag signifies lossless encoding.

To hear the full capability of Dolby TrueHD and DTS HD Master Audio, you must have a controller or receiver that can decode these formats, and connect the BD player with an HDMI cable.

In what seems a never-ending expansion of sound formats and the number of channels required to play them back, Dolby is making available a home-theater version of its Atmos surround-sound system, first developed for commercial cinemas. Atmos is unlike any other fixed-channel surround format. Rather than be restricted to a specified number of channels, the re-recording mixer can create a soundtrack with a virtually unlimited number of channels (128 in practice) for delivery to the theater. The same Atmos soundtrack can be reproduced by a varying number of loudspeakers in different theaters. The number of channels decoded and reproduced is determined by the number of channels in the particular movie theater. Atmos allows the re-recording mixer to specify within Atmos where in the theater each particular sound should be reproduced, regardless of how many speakers are installed in the theater. Metadata accompanying the soundtrack contain the instructions for where sounds should be positioned. The greater the number of speakers, the greater the spatial precision with which sounds can be rendered. Atmos can be adapted for 5.1-, 7.1-, and 9.1-channel home-theater playback systems.

12 Multichannel Audio

Introduction

Since the commercialization of the stereo LP in 1958, music has been delivered to consumers through two channels: left and right. Storing an acoustic event in two information channels was considered a compromise by early audio researchers (the literature suggests that some considered three channels a minimum), but no consumer format could hold more than two channels. An LP groove has only two sides, and multichannel tape machines were impractical for home playback.

The industry attempted to bring multichannel playback to the masses in the 1970s with Quadraphonic. As we all know, Quad was a commercial flop. The idea of a multichannel consumer-audio format died with Quad, and wasn't resurrected until the late 1990s, when a new generation of high-capacity, 120mm optical discs (SACD and DVD-Audio) was developed that could easily accommodate multichannel sound. SACD and DVD-Audio revived the idea of delivering to consumers more than two audio channels, but DVD-Audio is essentially dead, and SACD has become a niche format (although a sizable niche). The Blu-ray Disc format has a provision in its specification for an audio-only version that can store up to eight channels of uncompressed PCM audio at sampling rates of up to 192kHz and with 24-bit resolution. A few specialty labels have introduced BD music titles, but so far the major labels have shown no interest in the format, and audiophile labels have not rushed to reissue classic albums on BD. It's ironic that, just as a high-resolution multichannel format has become a reality, the major record companies, who control most of the world's music catalog, are concentrating on standard-resolution downloads.

In the past, quality-conscious consumers were stuck with compromised formats developed for the mass market. Creating a new physical format is a massive undertaking, and one that the major technology and music companies

embark on only if the financial rewards promise to be great. But the advent of streaming audio and downloadable music renders moot the idea of physical formats. We can, instead, bypass physical media in favor of lossless streaming and high-resolution multichannel downloads. The availability of highspeed Internet connections, coupled with ridiculously cheap hard-disk storage, makes it feasible to download high-res multichannel music. Nonetheless, if you're considering going multichannel, there are a number of factors to consider before making the leap.

Before we look at these new multichannel formats, discuss their merits, and talk about how to get multichannel audio into your home, let's take a brief look back at the history of multichannel audio.

A Short History of Multichannel Audio

Stereo audio recording and reproduction were invented by a young British engineer named Alan Dower Blumlein. At the age of 26, Blumlein was hired by the Columbia Gramophone company to develop a record-cutting system that didn't infringe on the patents held by Bell Telephone Laboratories. Although the cutting system he devised was monophonic, Blumlein was nonetheless thinking about a new technique that would capture audio signals with greater spatial fidelity. He was aware of the mechanisms by which humans locate the sources of sounds (phase and amplitude differences—see [Appendix A](#)), and was working in the background on a 2-channel recording-and-reproduction chain that would preserve and playback spatial positioning of sounds.

By early 1931, at the age of 29, Blumlein had conceptualized and built a system he called Binaural Sound—what we today call stereo. In that year he filed his landmark patent, “Improvements in and relating to Sound-transmission, Sound-recording and Sound-reproducing Systems” (patent number 394,325, filed December 14, 1931), which described in great detail a new microphone technique, a “shuffler” circuit, and record-cutting system, all of which made stereophonic sound possible. Blumlein described it this way:

The invention relates to the transmission, recording and reproduction of sound, being particularly directed to systems for recording and reproducing speech, music and other sound effects especially when associated with picture effects as in talking motion pictures. The fundamental object of the invention is to provide a sound recording and reproduction system whereby a true directional impression may be conveyed to a listener, thus improving the illusion that the sound is coming from the artist or other sound source as presented to the eyes.

Blumlein employed two bi-directional (also called figure-8 microphones, because their directional-sensitivity pattern looks like the numeral 8) crossed at 90°, with the two positive-lobe axes pointed at the center of the sound to be recorded. When the two mikes' signals are reproduced through two loudspeakers, the listener can localize sound between the two loudspeakers. (This explanation is simplified to remain within the scope of this book. For a detailed analysis, see John Eargle's *The Microphone Handbook*.)

Blumlein also invented other so-called “coincident” microphone techniques (M-S, X-Y) that are still used today, although only the crossed figure-8 technique is named after him. (Incidentally, you can hear “Blumlein” miking on many titles released by Water Lily Acoustics, notably the gorgeously recorded *A Meeting by the River* by Ry Cooder and V. M. Bhatt.)

Although Blumlein and his colleagues made many convincing recordings and films (of Blumlein himself walking across a stage while counting, of a train leaving a station, etc.), none of his contemporaries realized the significance of this work. The world would have to wait 27 years for the first commercial realization of stereo—the long-playing record, in 1958—long after Blumlein’s patent had expired. (Alan Blumlein was killed in a Royal Air Force plane crash in 1942, while working on radar research.)

About the same time, the Bell Telephone Laboratory was furiously researching all aspects of sound recording and transmission. In 1932, Harvey Fletcher of Bell Labs made a 2-channel recording with conductor Leopold Stokowski, but despite the obvious advantages of two channels over one, stereo recording and reproduction lay dormant until the late 1950s, on both sides of the Atlantic. Soon thereafter, of course, it became the standard for all recorded music.

The stereo record-cutting system developed by Blumlein, in which each channel is cut into one side of a 90° groove (the “45/45” system), appeared to limit the number of playback channels to two. In the 1970s, however, engineers developed several different techniques for encoding four audio channels within the two walls of a record groove. Although several competing systems were developed, the technique was generally called Quadraphonic. The two major Quad formats were SQ (developed by CBS, later bought by Sony) and CD-4 (developed by JVC). SQ was a matrix system in which the rear channels were matrix-encoded in the front channels. SQ suffered from extremely limited channel separation (as low as 3dB before logic steering became prevalent). CD-4 (Compatible Discrete 4) combined the front and rear channels in one signal, and recorded a difference signal between 30 and 50kHz. This high-frequency carrier was demodulated to obtain four discrete channels from a two-walled LP groove. CD-4 required a special phono cartridge and precise turntable setup. Other proposed Quad systems, such as QS and UMX, were never fully realized in the marketplace.

Quad was a musical and commercial disaster, primarily because the market was presented with competing systems and the demonstrations were laughably bad. In an attempt to show off the format’s 4-channel sound, Quad recordings often put the listener in the middle of a musical group, with instruments coming

from all directions. Moreover, the theory behind Quad was fundamentally flawed, as was repeatedly pointed out at the time by British audio genius Michael Gerzon. For one thing, Quad called for two loudspeaker pairs spaced at 90°—roughly one speaker in each room corner—which is too far apart to produce phantom images between the loudspeakers. The general public never cottoned to this gimmickry, and Quad quickly died. There are, however, Quad aficionados who still debate the relative merits of SQ and CD-4.

Although CD is a 2-channel medium (it has a provision for four discrete, lower-quality channels, but that capability was never used), it is possible to encode multichannel music on CD. Digital Theater Systems (DTS) has released a small number of multichannel titles encoded with its “perceptual coding” technique, which reduces the number of bits required to represent the audio signal. The result is some loss of fidelity (CD’s bitrate of 1.41 million bits per second for two channels must be spread over six channels), but a gain in the number of channels. Such discs must be played in a DTS-compatible player. Titles were also released in the Dolby Digital format, with the best examples coming from the classical label Delos. These formats were, however, stopgap measures until a true multichannel music carrier became a reality.

The availability of higher-capacity digital audio discs (SACD and Blu-ray), along with multichannel downloads, again offer the possibility of playing back music recorded with more than two channels of information. These formats obviate the serious technical compromises inherent in both SQ and CD-4 Quadraphonic, and don’t require a reduction in sound quality and lower bitrates to achieve multichannel capability. Consequently, we’re entering another age in which the general public will again be offered multichannel audio for music. The general public has, of course, embraced multichannel audio for film-soundtrack reproduction. The popularity of home theater has greatly increased the number of homes with a multichannel loudspeaker array (more than 60 million as of this writing).

Do We Want Multichannel Music Playback?

But just because we *can* play back multiple channels in the home, should we? Is multichannel audio a significant advance in sound quality, or a marketing gimmick foisted upon us? by greedy hardware manufacturers. This issue is perhaps the most polarizing controversy among audiophiles. Most audiophiles consider 2-channel music reproduction sacrosanct, and not something to be tampered with. Virtually all the recordings that we consider cultural treasures were recorded in stereo, and playing them back in any way other than exactly as they were intended to be heard is anathema. Moreover, surround-sound for music is associated with surround-sound for film, in which explosions and special effects are presented all around the listener—fine for movies, but justifiably abhorrent to the purist music lover. Moreover, 2-channel listening can be immensely satisfying, so why bother with surround sound? Multichannel audio is viewed by many audiophiles as a fad, not a legitimate means of advancing the goal of re-creating the original musical event in our homes with the greatest possible fidelity. On the other hand, some of the most ardent and dedicated audiophiles have championed multichannel as the most significant advance in audio since stereo.

In my view, multichannel audio is a quantum leap forward in advancing the music-listening experience—with a properly set-up system playing recordings made with musical sensitivity. There's no question in my mind that multichannel sound can greatly increase the spatial realism of reproduced music and deliver a more involving experience. The catch is that the system must be properly configured, and that the recording use the technology in a musically appropriate way rather than as an “effects” gimmick.

Multichannel overcomes a fundamental limitation of stereo, which is that stereo tries to create a 360° soundfield from two loudspeakers spaced 60° apart. Let's explore this idea with a thought experiment. Suppose we are in a concert hall in which we, but not the musicians, are enclosed in an acoustically isolated shell. If we cut two holes in the front of the shell facing the stage (analogous to listening to the performance through two loudspeakers), we'll hear mostly direct sound from the instruments. If we cut additional holes in the back of the shell, reflections from the back of the concert hall can enter the shell, and create a sound-field inside the shell that is closer to that in the concert hall. How could anyone argue that putting more than two holes in the shell is a step backward?

Reverberation and reflections arrive at our ears in the concert hall from behind us; why shouldn't they in the home?

I've attended many demonstrations of multichannel sound; I'll tell you about two of them that exemplify the problems and promise of multichannel music reproduction. In one demonstration, the sound from the speakers behind the listening position was played as loudly as the sound from the front speakers. I heard a trumpet blasting just behind my left ear, and an electric guitar behind my right ear. Most of the rest of the band was reproduced by the front channels, giving the impression not of hearing a musical group performing in front of me, but of being surrounded by the musicians. Not only was this unnatural and musically distracting, but hearing loud sounds suddenly blare out from behind me caused me to turn my head toward the sound source. Human beings have a survival instinct that makes them turn around when they hear a loud, sharp sound behind them. Putting the listener in an instinctive state of defensive readiness is not conducive to musical involvement. Moreover, hundreds of years of Western musical tradition tell us that the performers should be on a stage, and the audience in front of them. (Some compositions call for instruments behind the audience, but those are the rare exception.) To top it off, the overall sound was far too loud. All of this added up to a most unpleasant experience.

Now consider a second demonstration of surround sound I attended. For the original recording, five microphones had been placed in a large concert hall: three across the front, two in the rear. The five microphone signals were recorded in Sony's DSD format (the encoding format of SACD), and reproduced in the demonstration through five loudspeakers arranged similarly to the microphone placements. The re-creation of the original acoustic space was breathtaking; instead of hearing reverberation come from the front channels along with the instrumental images, the warmth of the hall was reproduced from behind me. More precisely, I wasn't aware of sound sources behind me, only of being inside a large acoustic space. Only ambience and hall reflections were reproduced by the rear channels, and at a *very* low level. Moreover, the front of the soundstage was more spacious, and the impression of "air" between instrumental images was greater than I've heard from any 2-channel playback. The soundstage didn't stop at the loudspeaker boundaries; it instead gently enveloped me, just as it does in live music. This microphone array and the multichannel delivery path had captured and reproduced the spatial aspects of musicians in a hall with far greater realism than is possible from two channels.

The promoters of the first demonstration wanted to hit people over the head with the fact that they could place instruments behind the listener. They had

absolutely no regard for whether the technology was being used in a musically meaningful way. In the second demonstration, surround-sound technology had increased the realism and provided a more satisfying experience than was possible with two channels. Unfortunately, the first example is the rule, the second the exception. Those who conducted the first demonstration apparently believe that most listeners are not sophisticated enough to discern and appreciate subtlety. It's demonstrations like the first one that will further polarize the audio community about the validity of multichannel audio.

This comparison makes the point that surround sound is neither a miracle technology nor the evil that many 2-channel purists consider it to be. Provided that it is used tastefully, multichannel audio has the potential of elevating the music-listening experience to a new level. If you have an antipathy toward multichannel audio, be sure your opposition is based on your own listening experience rather than on dogma, prejudice, or hearsay.

I'd like to tell you about another experience I've had with multichannel audio. Peter McGrath, one of the world's finest classical-music recording engineers, visited my listening room with his 4-channel Nagra D digital recorder and some 4-channel master tapes he'd made. Listening in multichannel mode for more than two hours, I was never conscious of sound originating from behind me. Rather, I felt enveloped in a large acoustic space. Halfway through a piece we switched to 2-channel, and the magic disappeared. The soundstage collapsed into the front channels, depth was reduced, and I suddenly had the feeling of hearing the music through a window in front of me rather than the illusion of being inside the concert hall. Had I heard the 2-channel recording first, I would have thought it excellent. But after hearing reverberation and hall ambience reproduced from behind me—the same place it originates with live orchestral music—there was no going back.

During this listening session, I had another epiphany. McGrath suggested that I walk around the listening room while the music was playing. Standing between the front speakers and facing the listening room's rear, I had the distinct impression of standing on a stage in a concert hall. The direct sound of instruments was to my left and right, and I heard reflections and reverberation in front of me coming from the rear speakers, just as one would when standing on a concert-hall stage. Standing between the rear speakers and facing forward, I heard more sound from the rear channels, which created the impression of being in the back of a concert hall. I could move forward a few steps and hear the spatial perspective change, exactly as I would in a hall, as the ratio of direct to reverberant sound changed. The effects of my movements were, of course,

exaggerated; a few steps in my listening room were equivalent to moving several dozen feet in a hall.

I then stood along the right sidewall, facing out into the room, with the front speakers to my right and the rear speakers to my left. Again, the impression of standing in that location in a concert hall—performers to my right, hall reverberation on the left—was eerily palpable.

This loudspeaker array created a continuous soundstage between the four loudspeakers. These soundstages are much easier to hear when moving around the room as I did; when facing forward, precise imaging between front and rear speakers was diminished because the sounds arrive from the side rather than from the front, which the brain processes differently from two front-arriving sounds.

This experience suggests that multichannel audio, done right, creates a more accurate, realistic, and compelling document of the original acoustic event. That the illusion survived my walking around the room tells me that the process is fundamentally valid. By contrast, the imaging we hear from 2-channel stereo is fragile: move your head sideways by a foot—or even a few inches—and the illusion can collapse.

Unfortunately, very few recording engineers currently view multichannel audio as an opportunity to better re-create the spatial aspects of live music. Rather, they see it as a new tool for fundamentally changing the listening experience. For example, AIX Records founder Mark Waldrep, a highly respected recording engineer and musicologist (he holds a Ph.D. in the subject), believes that multichannel audio is a means of escaping from the musical straitjacket of 2-channel stereo. His recordings are often made with the musicians in a circle around the microphones. The rear channels contain just as much information as the front channels—even with a string quartet recording. In my interview with him in Issue 134 (February/March, 2002) of *The Absolute Sound*, Waldrep articulated his approach: “For me, it’s an active choice to *not* create an archival document of a space where musicians are performing, and instead try to enhance or energize that performance with the technology we can bring to the equation.” Waldrep’s approach may be musically valid at some level—he knows far more about music than I ever will—yet I find such recordings disconcerting. To be fair, all AIX titles, which are superbly recorded at 96kHz with 24-bit resolution, contain both multichannel and 2-channel mixes, as well as “stage” and “audience” spatial perspectives—enough choices to satisfy any listener’s sensibilities.

At an Audio Engineering Society convention, I attended a workshop that explored how multichannel music should be recorded and mixed. The panel was a Who's Who of contemporary popular-music engineers. Each panelist took a turn presenting his views on how to use the additional channels. There was no consensus among the panelists; each engineer had a different approach. One engineer believed in placing the lead vocal *only* in the center channel. Someone else stated that the vocal should be mainly reproduced by the left and right channels, with slight augmentation by the center channel. Yet another advocated putting the vocal in the center channel, with slight augmentation in the left and right channels. And the discussion hadn't yet considered the rear channels.

You can judge for yourself the appropriateness of "hard" placement of the lead vocal in the center channel with James Taylor's *Hourglass* SACD. Play the recording through a multichannel speaker array and simply switch between the multichannel and 2-channel mixes. The multichannel mix puts Taylor's voice *only* in the center channel. The 2-channel mix puts his vocal equally in the left and right channels, which we hear as a phantom image between the two loudspeakers. In my experience with this recording, the hard placement of the vocal in the center channel kills the sense of depth and air around Taylor's voice.

At any rate, when it's used judiciously—however much experimentation that may require—multichannel audio can be a legitimate vehicle for musical expression, freeing the artist from the constraint of positioning sounds solely in front of the listener. In addition to this (potential) expansion of musical possibilities, multichannel audio can increase the listener's ability to hear deeper into the music. Psychoacoustic research has shown that humans can absorb more aural information when that information arrives from different locations. That is, a sound arriving from behind the listener is more intelligible than the same sound arriving from the front along with other sounds. We can better differentiate individual sounds, and thus more easily explore their musical meanings, when those sounds arrive from different locations around us.

This phenomenon was graphically demonstrated at a Consumer Electronics Show a few years ago. Tomlinson Holman, inventor of THX (which stands for Tomlinson Holman's eXperiment), set up and conducted the most unusual demonstration I've ever attended. Ushers with flashlights guided groups of showgoers to seats in a huge and nearly pitch-black room. The lack of light deprived us of any clue as to what we were about to hear—and prevented any prejudicial conclusions. Without any announcement describing what the demo was all about, the music began. The piece, an instrumental performed on synthesizer, was densely layered, yet perfectly intelligible on first listening. Each

of the many parts seemed to originate from a different point in space surrounding me. There were no discrete sound sources, only a seemingly continuous soundfield. I could easily shift my attention from one musical line to another. In fact, the unique sonic presentation encouraged—perhaps *compelled* is a more appropriate word—such an exploration of the music. It was an unforgettable experience.

After the demonstration, Holman turned on the lights and explained what we had just heard: a 10.4-channel audio system playing a piece composed and performed by Herbie Hancock that had been recorded in 10.4 discrete channels (ten full-range channels and four subwoofer channels). Hancock had used this wide sonic palette to expand his musical expression. My enthusiasm for what I'd just heard was tempered by the fact that a 10.4-channel audio system is never likely to become a commercial reality. Still, Holman had proved his points: multichannel audio opens up new musical horizons, and more channels are better than fewer channels.

An unfortunate side effect of the home-theater boom is that multichannel music formats and playback systems have been dictated by the movie industry, rather than developed from first principles for the highest possible musical fidelity. Today's multichannel loudspeaker layouts, channel configurations, and use of a single subwoofer all trace their origins to the film-sound industry. For example, a center loudspeaker is required in movie theaters because of the very wide placement of the left and right speakers. Dialogue must seem to originate onscreen; in fact, the center channel in movie sound is called the "dialogue channel." For music reproduction, we may want a center *loudspeaker*, but not a center *channel* to drive it. (This is explained later in the discussion of Trifield processing.)

A fundamental problem of judging multichannel audio is the scarcity of good setups and demonstrations. Most multichannel loudspeaker arrays at high-end audio dealers are optimized for home-theater playback, not music listening. After all, home theater is a booming industry; multichannel audio is not.

Ironically, the best demonstrations of multichannel audio's capabilities may be in the car. The big advantage of a car stereo is that the loudspeaker placements, listener position, and acoustic environment are all known. There's no chance for a botched demonstration because of less-than-ideal speaker placement or other variables. Of course, a car audio system, no matter how good, can't approach the quality of a true high-end system set up in a room. Nonetheless, the multichannel audio system and overall high quality available in a few automotive models offer the general public a good representation of multichannel audio's potential.

How to Get Multichannel Audio in Your Home

If you decide to take the multichannel plunge, you can follow one of four paths:

- 1) Build a system optimized for 2-channel music playback and home-theater reproduction, and accept compromised multichannel music performance.
- 2) Build a system optimized for 2-channel and multichannel music playback, and accept compromised home-theater sound.
- 3) Build a system optimized for 2-channel and multichannel audio without any capability for reproducing film soundtracks.
- 4) Build a system that does it all—that’s equally adept at the playback of 2-channel and multichannel music and of film soundtracks.

This first approach is the most common, which is why it’s difficult to hear state-of-the-art multichannel audio demonstrations. Such a system employs sidewall-mounted dipolar loudspeakers, which are ideal for film-soundtrack reproduction but not for multichannel music. As described in [Chapter 11](#), dipolar surround speakers mounted on the sidewalls better create a sense of envelopment and diffusion in the surround soundfield. They “smear” the rear-channel information by their directional pattern; the listener hears their sound only after the sound has been reflected around the room. Although this smearing may be appropriate for the atmospheric sounds in a film soundtrack (a rainstorm, jungle scene, or cityscape, for examples), when applied to a multichannel music mix this smearing is a distortion of the musical presentation created by the recording engineer. Multichannel music is recorded and mixed with unipolar loudspeakers, and should be played back with an identical loudspeaker array. You will never hear aural images between dipolar surround loudspeakers.

The second option—optimize your system for multichannel music at the expense of compromised film-soundtrack performance—simply involves using unipolar rather than dipolar surround loudspeakers. These need to be positioned behind the listener rather than on the sidewalls, as would dipolar models. When listening to multichannel music, you’ll hear a better representation of what the recording engineer intended, but film soundtracks won’t sound quite as spacious or enveloping. The atmospheric sounds in the surround channels will appear to

originate from two points in space (the left and right surround speakers) rather than envelope you in a diffuse wash of sound. Nonetheless, many listeners find unipolar surround speakers perfectly acceptable—even preferable—for film-soundtrack reproduction.

Both of these options require a home-theater controller (described in [Chapter 11](#)), which will function as a surround-sound decoder for film soundtracks, and as a source selector and volume control for multichannel music. The six analog outputs from a universal-disc player feed the controller's multichannel analog input. Alternately, the controller can accept digital multichannel audio from a Blu-ray Disc player (or universal disc player) via an HDMI cable. The digital outputs from your other sources (satellite, DVD player) feed the controller's digital inputs.

Which brings us to the third option: dismissing film-soundtrack reproduction entirely. If you never intend to use your system to watch movies, you'll obviously choose unipolar surround loudspeakers. Although you won't necessarily need a home-theater controller, you will need a way of selecting sources (one of which will be multichannel) and controlling the volume. The solution for this system is the multichannel preamp. This new breed of product includes no provisions for surround-sound decoding, bass management, or other home-theater requirements. Rather, it is simply a preamplifier with six channels rather than two. [Fig.12-1](#) shows an example of a multichannel preamplifier. Note. at the lower left, the two sets of multichannel inputs, balanced and unbalanced.

Finally, if you want the ultimate performance for both multichannel music and movies, you'll install two sets of surround speakers: dipolars on the sidewalls and unipolars toward the rear. You'll also need two separate 2-channel power amplifiers, one to drive the dipoles, and one to drive the unipoles. A pair of Y adapters on the controller's surround-channel outputs splits the signals so that the surround signal can drive the two amplifiers. You select which surround loudspeakers are engaged by which power amplifier you turn on. This method is a bit cumbersome, but it does solve the music/movie dilemma.

All these approaches require a multichannel source such as a universal or Blu-ray player, or a music server with multichannel capability. Such machines are described in detail in [Chapters 7](#).



Fig.12-1 A multichannel preamplifier—rather than a digital controller—provides the highest audio quality for multichannel music reproduction. (Courtesy Audio Research Corporation)

Loudspeaker Types and Placement

The home-theater paradigm calls for five (or seven) loudspeakers and a single subwoofer. The subwoofer is driven by the low-frequency effects (LFE) channel in a Dolby Digital soundtrack and, selectively, bass from any of the other channels. In fact, many home-theater systems employ small satellite loudspeakers with a single subwoofer that reproduces all of the bass. Even if your home-theater speaker system has full-range left and right speakers, chances are that the center loudspeaker will be the small, horizontally mounted variety.

This arrangement is, however, not ideal for multichannel music reproduction. The best speaker configuration for multichannel music is five identical full-range speakers in the locations shown in [Fig.12-2](#). No subwoofer is required. Of course, if your system also functions as a home theater, you can have a subwoofer that is engaged only for film-soundtrack reproduction, not music playback. The loudspeakers should be equidistant from the listener. This array creates a what is apparently a single continuous soundstage, but is actually four separate soundstages: between the left and right loudspeakers; between the surround loudspeakers, between the front left and rear left loudspeakers; and between the front right and rear right loudspeakers.

Chesky Records objected to the film-sound industry dictating music playback standards, so the company developed its own channel configuration and loudspeaker array, and made some recordings for playback on what it considered its ideal system. That system, called Chesky 6.0, uses the loudspeaker array shown in [Fig.12-3](#).

Demonstrations of this technique have sounded exceptionally good. Unfortunately, it's unlikely that this excellent approach will become a standard.

Bass Management

Bass management is a subsystem within a multichannel source machine (BD player, SACD player) or digital controller that allows you to selectively distribute bass information among your speaker array. For example, if you have full-range left and right loudspeakers but small center and surround speakers, bass management keeps bass out of the small center and surround speakers and directs that bass to a subwoofer (if the system includes one). Bass management is essential in a multichannel speaker system because not all systems use full-range speakers.




Fig.12-2 Loudspeaker placement for multichannel music reproduction




Fig.12-3 Loudspeaker placements for the Chesky 6.0 technique. (Courtesy Chesky Records)

Although bass management is provided in both multichannel players and controllers, you need set up only one device to perform bass management. Note that if you choose a system with a multichannel preamp rather than a digital controller, you'll need to set up bass management in the player.

Calibration

Just as it's vital to 2-channel playback that the left and right loudspeakers produce identical sound-pressure levels, the speakers in a multichannel audio system need to be precisely calibrated so that they all produce the same level at the listening position. Adjusting the individual channel levels compensates for differences in loudspeaker sensitivity, placement, distance from the listener, and the effects of room acoustics.

All digital controllers are equipped with a pink-noise generator that can be programmed to cycle around all the channels. Pink noise is a hissing-like sound. You'll need to use this test signal in conjunction with a sound-pressure level (SPL) meter; setting the levels by ear is not accurate enough. Available from RadioShack for about \$30, an SPL meter lets you accurately calibrate your multichannel audio system. RadioShack offers digital and analog SPL meters; buy the easier-to-read analog type. Switch the meter to the C-Weighted position, Slow Response, and set the knob to 70. Position the SPL meter where your head is when you sit in the listening seat, and point the meter toward the ceiling. (This ensures that the meter doesn't favor anyone channel.) Turn on your controller's test noise, and set each channel's volume until the meter's display reads 5. This indicates that the noise is being reproduced at a level of 75dB. Repeat this procedure for each channel. You may need to repeat this cycle for all the channels a few times to get the levels precisely matched.

After calibration, the controller's volume display will accurately show the level relative to 0dB. A typical listening level may be displayed as "-10dB. This may be confusing; why is 0dB the reference level? As explained in [Appendix A](#), decibels express a difference, or ratio, between two values rather than an absolute value. At the 0dB volume setting, the system should produce a peak SPL of 105dB. This calibration standard comes from film sound, where the maximum SPL of the loudest peaks is 105dB. All other volume settings are referenced to this maximum value. With the system properly calibrated and the controller set to 0dB, the SPL in your listening room should be the same as on the film-sound mixing stage.

If the surround channels are too high in level, you'll hear too much ambience and reverberation and not enough direct sound. Set the rear channels too low in level and you won't experience the full surround effect. It's a little like changing rows in a concert hall: the farther from the stage you sit, the lower the ratio of direct to reflected sound. Unfortunately, multichannel music discs in which the surround channels contain reverberant information vary greatly in the amount of

signal in the surround channels. This means that one calibration is not ideal for all discs. It's best to start with the precise meter-derived calibration, and adjust slightly from there to suit the recording.

You'll also need to set the controller's individual channel delays so that the sounds from all of your speakers arrive at the listening position at the same time. The controller's setup menu will ask you for the distance from each loudspeaker and to the listening seat. The controller then calculates the amount of delay to apply to each channel. If the center channel has too little delay, soundstage width will be constricted, with instruments and voices tending to bunch up in the middle. The middle of the soundstage will seem more forward than the sides. Conversely, too much center-channel delay will increase the sense of depth, but at the expense of presence and tonal fidelity.

Multichannel Playback from 2-Channel Sources

The proliferation of multichannel loudspeaker arrays, combined with the scarcity of multichannel recordings, has led to the creation of decoding algorithms that generate 5- or 7-channel playback from unencoded stereo program material. Consumers who paid for those center and surround loudspeakers want to hear them working all the time, not just when watching movies. That somewhat cynical view aside, unencoded stereo programs often contain ambience information that is sufficiently decorrelated from the direct component of the signal that this ambience can be extracted and reproduced by rear loudspeakers to “unlock” the recording’s spaciousness—sometimes with stunning effect.

Among these processes are Dolby Pro Logic II, Pro Logic IIx, DTS Neo:6, Logic 7, and Music Logic. All create signals to drive the center and surround loudspeakers from 2-channel sources. These decoding circuits are standard in virtually all digital controllers, except for Logic 7, which is available only in products made by Harman International (Harman/Kardon, Lexicon, JBL, for examples). Another example is EARS, developed by NAD Electronics. EARS is relatively gentle, feeding in-phase information common to the left and right channels to the center speaker, and extracting ambience that feeds the rear loudspeakers. The effectiveness of these systems varies greatly with the recording; some sound worse with processing, others sound the same (if there is very little uncorrelated ambient information), and a few sound spectacular—almost like a discrete multichannel recording.

These decoding algorithms all attempt to extract ambience *already present in the recording* and deliver it to the rear loudspeakers. These techniques are not to be confused with the digital-signal-processing (DSP) modes in home-theater receivers and controllers with labels such as Concert Hall, Jazz Club, or Stadium. Rather than extract ambience from the original program, those DSP modes *synthesize* reverberation and add it to the music. Attempts have been made to measure the acoustic properties of concert halls and other venues and encode that acoustic signature in DSP algorithms that can be overlaid on the musical signal. All these schemes are nothing more than marketing gimmicks; still, many consumers choose A/V receivers based on the number and variety of these surround modes. (One A/V receiver I know of offers 1024 possible surround-mode variations.) Needless to say, such manipulation is antithetical to high-end values, and should be avoided. (Some would argue that even ambience recovery-

techniques that only extract information already in the recording are antithetical to high-end values.)

I've found the most effective and musically appropriate algorithm for playing back 2-channel sources over more than two loudspeakers to be Trifield, a process found in products made by Meridian Audio. According to Meridian, "Trifield extracts the mono and surround components of the original recording. It then calculates the signals for the front left, center, and right speakers, using the phase and amplitude differences between the three front channels, to redistribute the sounds on a frequency-dependent basis. This gives a significant improvement over traditional stereo, which converts the differences between the microphone signals into amplitude differences in the speaker signals. Trifield is recommended for well-made recordings and stereo television broadcasts that are not Dolby Surround encoded."

Trifield, which is derived from Ambisonics (described below), actually comprises two algorithms. The listener has the choice of frontal stereo only, or the addition of extracted ambience in the rear channels. Meridian controllers allow you to select side or rear surround loudspeakers, or both. The extracted ambience signal is extremely subtle; the listener is never consciously aware of sound sources behind or to the side. Rather, the extracted ambience seems to make the front soundstage more spacious, open, and detached from the speakers. A Width control allows you to widen the soundstage. A wider soundstage is possible without "hole-in-the-middle" imaging because the center speaker anchors the center of the soundstage.

In my experience, listening in Trifield does all of that: it fills in the center of the soundstage, presents a wider sweet spot, and creates more stable images between the speakers. Moreover, the imaging doesn't shift with small head movements, as it often does when listening to two loudspeakers, the so-called "head-in-a-vice" syndrome.

Here's an interesting story about Trifield: During a visit to my listening room by an electronics designer who was someone who believes that music should be reproduced by just two loudspeakers, I performed an experiment. We listened to music for about an hour, but unbeknownst to him, I had Trifield engaged. Part of the way through a piece of music, I switched to 2-channel stereo. He was instantly and visibly let down by the sudden change. He heard a reduction in sound quality, but had no idea what could have caused it—until I gave away the secret. Trifield is subtle—it's not an obvious "effect," but profound in its ability to solidify the soundstage.

It's important to keep in mind that there's nothing sacrosanct about using two loudspeakers to reproduce music. Early researchers considered two channels a compromise dictated by practicality and the technology of the day. Further, the word *stereo* doesn't mean "2-channel" reproduction. Rather, it comes from the Greek word for *solid* because two channels can create apparently solid, or three-dimensional, images between the speakers.

Ambisonics

The solution to how best to capture a soundfield, and to reproduce that soundfield through more than two loudspeakers has been solved for more than 40 years. It's called Ambisonics. Invented by Michael Gerzon, Ambisonics is a simple expansion of Alan Blumlein's idea, adding an omnidirectional microphone to the crossed figure-8 microphones that constitute a Blumlein pair. This setup captures all the information (in the horizontal plane) present in the original soundfield. When reproduced by four or more speakers (the more speakers, the better), Ambisonics reconstructs the original acoustic soundfield. This isn't merely a marketing claim; there's a solid body of mathematics that shows that Ambisonics indeed encodes and reconstructs every aspect of the soundfield.

An additional pair of figure-8 microphone can be added to capture the height information as well. In Ambisonics recording, the monaural sound pressure is called the "W" component, the left-right velocity component is called "Y," the front-back component is called "X," and the optional up-down (height) component is called "Z."

The W, X, and Y signals are called the B-Format and are obviously not compatible with 2-channel distribution formats. The signals can, however, be combined to produce stereo (and mono) compatibility. This encoding method, called UHJ, matrix-encodes the three (or four, with the Z channel) signals for mono, stereo, and multichannel compatibility. UHJ-encoded material is decoded on playback by an Ambisonics decoder in a digital controller.

Ambisonics recordings are made without regard for the number of speakers through which the program will be reproduced. The special microphones sample the sound-field and encodes all the necessary information about that soundfield. On playback, the Ambisonics decoder generates the appropriate output signals based on the number of speakers in the system. Because the multichannel audio is hierarchically represented, it can be played back through as few or as many speakers as your room can accommodate. The B-Format representation is completely independent of the loudspeaker array; the Ambisonics decoder will correctly map the soundfield to the number of loudspeakers.

Many hundreds of Ambisonics recordings have been made in the past 40 years, notably by the classical label Nimbus, which uses the technique exclusively,

and released an extensive series of Ambisonics recordings on DVD-Audio in 2004. Those recordings may be available for high-multichannel download in the future. Nimbus at one time owned the Ambisonics trademarks, but sold them to Meridian Audio. Even if the signal will never be decoded, Ambisonics has many advantages for 2-channel recording and reproduction.

The advent of high-capacity, multichannel digital-audio discs such as Blu-ray provides an opportunity to deliver to the listener the full Ambisonics experience. It is possible to encode the hierarchical Ambisonics information in a multichannel signal and flag that data so it can be reproduced by a 5.1-channel speaker system, even if the playback system doesn't include an Ambisonics decoder. A hierarchically encoded multichannel signal is not tied to any particular number of speaker feeds; rather, it is a method of storing and transmitting multichannel audio to fit a wide variety of speaker configurations. The hierarchical signal is "scalable" to the number of loudspeakers through which the signal will be reproduced, thus making a multichannel signal perfectly compatible with any number of playback channels, including mono.

Ambisonics was developed long before any standards were in place that specified the number of loudspeakers in homes. However, today's ubiquity of 5.1-channel playback systems presents an opportunity to deliver Ambisonics without a decoder in the playback chain. With the number of speakers known in advance, the Ambisonics signal can be "pre-decoded" for playback through a 5.1-channel speaker array. Although no decoder is required, it isn't possible to use *more* than five loudspeakers with this so-called G-Format. In addition, height information in the G-Format can't be reconstructed without a decoder. If the system does employ an Ambisonics decoder, that decoder can be programmed to generate the appropriate number of signals to feed the speaker system. This decoding is automatically engaged by the datastream flag, making the process completely transparent to the user.

The ideal speaker arrays for Ambisonic s playback are four speakers in a rectangle; five speakers in a regular pentagon; or six speakers in a regular hexagon. As mentioned earlier, the greater the number of speakers, the more faithful the reconstruction of the original soundfield that was sampled by the microphone array.

There's no question that Ambisonics is by far the most musically appropriate multichannel recording technology. It reconstructs the soundfield in a way that 5.1-channel systems can't approach, producing a much more believable sense of space, more precise image placement, and a much larger listening area than the typically small, sharply limited "sweet spot."

So why didn't Ambisonics catch on? First, Ambisonics came to market just as Quad was dying. Ambisonics was lumped in with Quad's failures, and record companies and electronics manufacturers wanted to put that failure behind them. Ambisonics was also seen, mistakenly, as just another competing Quad format. Moreover, no major record labels had any interest in Ambisonics, because all had already associated themselves with one of the Quad formats. Ambisonics was also seen as an audiophile technique suitable only for capturing classical music in a concert hall with natural ambience, and not appropriate for studio records (although Alan Parsons's Ambisonics-encoded *Stereotomy* destroyed that argument). Finally, Ambisonics was doomed by a complete marketing failure, the result of the patents being turned over to a British government agency whose job was to promote and market British inventions. The agency was ill-equipped to promote Ambisonics and negotiate licensing agreements with hardware manufacturers.

Still, many Ambisonics titles are readily available today, and Meridian Audio incorporates Ambisonics decoders in all of its digital controllers. It is possible to encode Ambisonics hierarchical data in multichannel datastreams, which may breathe new life into the technology. Just as it took 28 years for Alan Blumlein's invention of stereo to become a widespread commercial reality, Ambisonics may one day become the dominant multichannel music format.

More information about Ambisonics, including a comprehensive list of titles and decoding hardware, can be found at ambisonics.net.

13 Personal Audio: Headphones, Desktop, and Mobile Audio

Introduction

The personal-audio revolution began in on July 1, 1979, when Sony launched the Walkman, a portable cassette-tape player. Although clunky by today's standards, the Walkman liberated music from the home and car, letting listeners take music anywhere. The Walkman was a massive commercial success—Sony sold more than 200 million of them, and that's not counting the hundreds of millions of portable cassette players offered by myriad other companies.

Personal audio was reinvented on October 3, 2001, by Apple's pocket-size stereo music player, the iPod. To date, consumers have purchased more than 350 million iPods. But even the iPod's days are numbered, as music playback is increasingly incorporated into every possible mobile electronic device, from tablets to smartphones. These devices can also stream music from the Internet through services such as Tidal, giving you virtually unlimited access to any music at any location.

Coincident with the iPod's rise, the market for headphones exploded in one of the biggest product-category booms in the history of audio. Clearly, people like to listen to music anywhere and everywhere, but headphones aren't just for listening on the go; exotic high-end designs, and specialized amplifiers to drive them, have provided exceptional sound quality for a fraction of the price of a traditional loudspeaker-based system. Moreover, a headphone-based system doesn't require giving up a significant portion of your living space to a hi-fi—and it's much easier to move when that time comes.

The advent of computer audio has also made listening at your desk a natural; you already have the hardware for a music server (your computer) right in front of you—all you need is a few ancillary products to enjoy music on your desktop.

In this chapter I look at all aspects of personal audio, including headphones, headphone amplifiers, desktop systems, and portable music players.

Headphones

The broad term *headphones* is often used to describe any personal listening device attached to the head. But the category encompasses four distinct subcategories: earbuds, earphones, in-ear monitors, and headphones. Let's look each of these in turn.

Earbuds, Earphones, and In-Ear Monitors

Earbuds are the ubiquitous devices supplied with most portable music players and mobile devices. Earbuds have poor sound quality, which is one reason headphone sales have soared along with the availability of music-playing mobile devices: millions of listeners have discarded their earbuds and bought better-sounding earphones, in-ear monitors, and headphones.

Earbuds are nonstarters for high-quality audio, so I'll move on to *earphones*. There's a single difference between earbuds and earphones: earbuds fit in the outer ear, but an earphone is inserted in the ear canal, where, ideally, it creates an airtight seal with the canal. The difference in sound quality conferred by this apparently small distinction is profound. Creating an airtight seal greatly improves sound quality, particularly the bass response. A tight seal also prevents outside noise from intruding on the music. Earphones are often supplied with a variety of interchangeable eartips, made of soft silicone or compressible foam, which allows you to choose the tips that provide the best fit with your ears. Although earphones may be superficially similar to earbuds, very-high-quality sound is possible from good earphones, never from earbuds. Fig.13-1 shows a high-end earphone.



Fig.13-1 Earphones offer much better sound quality than earbuds. (Courtesy Cardas Audio)

Although good earphones can provide superb sound, how well those earphones fit your ear canals—and thus how good they sound—is somewhat hit or miss, even with the range of flexible eartips supplied with the earphones. Because earphones' sound quality depends so much on how well they fit the ears, some manufacturers offer *in-ear monitors*—earphones that have been custom molded to your ear canals.

A little history of how in-ear monitors came about helps in understanding their *raison d'être*. Remember seeing, at amplified concerts, those small black speakers sitting at the front of the stage, angled up toward the musicians? Those were stage monitors, which allowed the musicians to hear themselves playing and singing and have largely been replaced by in-ear monitors. Stage monitors are problematic—they're bulky, tend to generate feedback, must be turned up very loud to be heard above the house mix, and require that each musician stand in front of his or her individual monitor. The professional live-sound industry solved these problems with the custom-fit in-ear monitor. Giving each musician personal in-ear monitors made stage monitors a thing of the past. But for in-ear monitors to be effective, they must form a perfect seal between the transducer (the tiny speaker) and the ear canal. Any leakage is unacceptable, because the sound from the huge main speakers would swamp the sound from the in-ear monitors and confuse the musician.

The solution is to create impressions of a person's ear canals, then mold each earphone of a pair to provide a perfect fit with the corresponding canal. High-end audio has adapted this approach to creating very-high-quality in-ear monitors for music listening at home or on the go. The in-ear monitor not only provides high isolation from noise; the airtight fit allows the tiny transducer to perform optimally. A high-quality in-ear monitor can deliver spectacular performance, including bass of a depth and power not normally associated with tiny personal listening devices. Another benefit is comfort. Because in-ear monitors precisely match the shapes of your ear canals, you can almost forget they're there. Moreover, the slightly larger size of in-ear monitors relative to earphones gives designers room to include more elaborate transducers. Finally, a custom-fit in-ear monitor can provide as much attenuation of background noise as a headphone employing active noise-canceling technology (described later in this chapter).

To create the molds from which in-ear monitors will be cast, an audiologist trained in creating in-ear impressions injects a semiliquid silicone or acrylic material into your ear. During this process, you'll be asked to hold a plug in your opened mouth, to prevent your ear canals from moving as the material solidifies. About 15 minutes later, the audiologist removes the now-solid material, which retains a perfect impression of your ear canal. Each in-ear monitor is hand-cast from such a mold to create a shape unique to your ear. Note that the entire in-ear monitor—not just the tip—conforms to your ear canal's shape. To ensure an absolutely perfect fit, some in-ear monitors are made from a pliable material that softens slightly when warmed by your ear canal. Many makers of in-ear monitors

bring an audiologist to consumer hi-fi shows to create your ear-canal impressions right on the spot. The finished monitors are shipped to you several weeks later.

Because the market for musicians' in-ear monitors is relatively large and the customers are so demanding, much research has gone into designing good-sounding products. Some designs are extremely elaborate, and may include five drivers and two crossover points—one bass driver, two midrange drivers, and two tweeters, for example. Others rely on a single full-range driver. No matter the number of drivers, they are usually miniature versions of the moving-coil drivers found in freestanding loudspeakers. Another type of transducer used in in-ear monitors is called the *balanced-armature driver*, in which the armature is a tiny rod balanced on a pivot point, like a teeter-totter, within a magnetic field. The audio signal is applied to a coil surrounding the armature, which creates a fluctuating magnetic field that magnetizes the armature. The armature's magnetic field reacts with the fixed magnetic field created by the magnet, causing the armature to oscillate and, with it, the tiny diaphragm connected to one end of the armature. An upper-end in-ear monitor may contain several balanced-armature transducers, each optimized for reproducing a different range of frequencies. Balanced-armature drivers are extremely efficient, requiring very little power to produce sound.

Custom-fit in-ear monitors require a greater commitment of time and money than off-the-shelf earphones, but their combination of comfort, noise isolation, and, most important, sound quality, are unequaled among in-ear listening devices.

Headphones

Many listeners prefer headphones to in-ear products, particularly those for whom portability isn't important. As good as high-quality custom in-ear monitors can sound, it's headphones that can provide a state-of-the-art personal listening experience.

Which are the best headphones? There's no single answer to this question, but as many right answers as there are listeners and specific applications. Choosing the right headphones for you begins with defining how you will use the headphones, what combinations of sound and comfort you most value, and your budget. Let's explore some of these criteria, keeping in mind that many of them overlap.

Location and Intended Use: Where will you use the headphones? If you plan to listen in public places such as airplanes, trains, or subways, you'll want the headphones to provide isolation from outside noise as well as prevent those nearby from hearing your music. Earphones and in-ear monitors are ideal for this, both for their portability and their noise-isolation qualities. Earphones that provide a good fit—and *all* custom in-ear monitors—provide excellent isolation from noise. But remember that keeping outside noise from intruding on your listening experience can also prevent you from hearing sounds that warn of danger.

Conversely, if you're buying headphones because your home system is headphone-rather than loudspeaker-based, or to enjoy music late at night without disturbing neighbors or family members, noise isolation and portability won't be as important. For home listening, your top priorities should be sound quality and comfort.

Supra-aural and Circumaural: Headphones are classified by how they fit on your head: supra-aural and circumaural. *Supra-aural* (literally, "over the ear") headphones have pads that rest on the ears (Fig.13-2), which is why they're also called "on-ear" headphones. By contrast, a *circumaural* (literally, "around the ear") headphone completely covers and encloses the entire ear (Fig.13-3). Supra-aural headphones are generally lighter, smaller, and more easily fit in a travel bag. The disadvantage is that they allow in outside sounds, which can be distracting. And it works in the other direction: Supra-aural headphones leak sound from the headphones to the outside world, potentially disturbing those nearby. Supra-aural

headphones are thus not a good choice for subway commuters or airline passengers. By completely enclosing the ear, circumaural headphones provide greater isolation from outside sounds, and those nearby won't hear your music. If you plan to listen in the presence of people who may be bothered by your music, choose circumaural over supra-aural headphones.



Fig.13-2 Supra-aural headphones rest on your ears, and are also called “on-ear” headphones. (Courtesy Bowers & Wilkins)



Fig.13-3 Circumaural headphones surround the entire ear. (Courtesy Beyerdynamic)

Open-Back vs. Closed-Back: A second major division in headphone design is open-back and closed-back headphones. In open-back 'phones ([Fig.13-4](#)), the

side of the diaphragm facing away from the ear is open to the outside world. Conversely, a closed-back design (Fig.13-5) completely seals the driver in an enclosure. Open-back headphones provide less isolation from outside noise, and allow others nearby to hear your music.

Generally speaking, open-back headphones have a more spacious soundstage and better imaging, and closed-back headphones have deeper, more powerful bass. Completely enclosing the driver, however, causes the sound radiated from the back of the diaphragm to be reflected from the enclosure back to the diaphragm, introducing unwanted diaphragm motions that are heard as colorations and smearing. Generally, open-back 'phones offer better sound quality than closed-back designs. Closed-back 'phones can also make your ears feel hotter with extended use; open-back 'phones “breathe.”



Fig.13-4 Open-back headphones often sound better than the closed-back variety, but offer less isolation from outside sounds. (Courtesy Grado Labs)



Fig.13-5 Closed-back headphones are the best choice when you don't want to disturb those around you. (Courtesy PSB Speakers)

Portability: The ultimate in portability is provided by earphones and in-ear monitors, which are often supplied with a small travel case that easily fits inside a briefcase, backpack, or purse. Some listeners object to inserting objects in their ears, and will opt for headphones even for traveling. If portability is important, look for headphones that fold up for storage in a travel case.

Comfort: The best-sounding headphones in the world won't be much good if they're uncomfortable to wear. Headphones vary immensely in how they feel on the head and against or around the ears, and any discomfort will only grow the longer you wear them. There's some agreement about which model headphones are generally comfortable and which aren't, but everyone's head is different. You should wear the headphones under consideration for an extended period, if possible, before committing to a purchase.

Generally, lighter headphones are more comfortable than heavier models. Supra-aural headphones that rest against the ear usually weigh less than circumaural 'phones that enclose the ear, but the pressure applied to your ears by supra-aural headphones can become uncomfortable. Also consider the headband padding, which ranges from a vinyl cover over bare metal to a deep cushion. Because every listener's head has a different shape and size, some brands or models of headphones will naturally fit you better than others.

Ruggedness: If your headphones will never leave your listening room, rugged build quality won't be an important consideration. For those who carry their

'phones everywhere they go, choosing headphones designed to withstand the rigors of daily travel is essential to realizing a long service life. Headphones designed for professional use are usually more durable.

Sound Quality: The criteria for judging the sound quality of headphones and loudspeakers are the same: You want a smooth tonal balance with no colorations, particularly through the midrange. The bass should be extended and full, but not bloated and thick. Headphones vary considerably in their treble extension and sense of openness, and a model without good treble extension will sound closed-in. Transparency, resolution, and dynamic range are all important factors in how much musical satisfaction the headphones will deliver in the long term. Listen to familiar music through the headphones you're considering, preferably for an extended audition.

Noise Canceling: Active noise canceling is a technology that suppresses steady-state background noise, such as that generated by airplanes, cars, and trains. A microphone picks up ambient sound, amplifies the sound, inverts its polarity, and drives the headphones with the inverted signal. This inverted signal cancels, to some degree, the ambient sound leaking into your ears through the headphones. Noise-canceling headphones greatly reduce fatigue on long flights. By lowering the level of background noise, this technology can also make the music more intelligible and offer greater resolution.

Noise-canceling headphones require a power source: either an integral rechargeable battery pack, or two or more AA or AAA replaceable batteries. They're a little larger and more expensive than conventional headphones, but if you travel frequently, they can go a long way toward making travel less unpleasant. There's no reason why noise-canceling headphones need to compromise fidelity; I own a pair of terrific-sounding headphones that happen to have noise-canceling technology. They greatly reduce the constant roar inside an airplane cabin; I wear them even when I'm not listening to music.

Amplifier Matching: As explained in technical detail later in this chapter, headphones perform best when their electrical characteristics match those of the amplifier driving them. The primary specification to look at is the headphones' impedance. Briefly, inexpensive portable audio devices tend to work best when driving headphones of high impedance (i.e., higher than about 100 ohms). Low-impedance headphones perform better when driven by a more robust amplifier,

such as those found in portable DACs and dedicated headphone amplifiers. See “Headphone Specifications,” later in this chapter, for more detail on this subject.

Balanced and Unbalanced Drive: Some headphones offer the option of balanced connection between the amplifier and headphones. Balanced connection, described in detail in [Chapter 10](#) (“Cables and Interconnects”), carries the audio signal on three conductors rather than two. Balanced connection requires headphones that are wired with balanced connectors, as well as an amplifier designed for balanced operation, and is described in greater detail later in this chapter, under “Headphone Amplifiers.”

Wireless Headphones: Wireless headphones free the user from being tethered to the amplifier, but that convenience comes at a price: sound quality. Wireless headphones often use Bluetooth for transmitting the signal between the base station and the headphones, compromising fidelity. Bluetooth employs a coding and decoding system to reduce the bitrate and thus the wireless signal’s bandwidth. For that reason, wireless headphones aren’t recommended except for casual use, or with the multichannel headphones described later in this chapter.

Dynamic, Electrostatic, and Planar-Magnetic Headphones

As we saw in [Chapter 6](#) (“Loudspeakers”), an electrical signal can be converted into sound by various technologies, including the dynamic moving-coil driver, the electrostatic panel, and the planar-magnetic transducer. You’ll find a similar array of technologies among highend headphones. Each of these headphone technologies is simply a scaled-down version of the dynamic, electrostatic, and planar-magnetic drivers found in freestanding loudspeakers. (A complete explanation of how each of these technologies works is included in [Chapter 6](#).)

The vast majority of headphones have dynamic drivers. *Dynamic* headphones are rugged, small in size, and low in weight relative to electrostatic and planar-magnetic headphones. Although some dynamic headphones are expensive, they are generally less costly than electrostatic and planar-magnetic designs. The headphones shown in [Figs. 13-2](#) and [13-3](#) are dynamic designs.

Electrostatic headphones are rare, but they have a cult following for their outstanding transparency, resolution, and speed. These qualities, which are shared with electrostatic freestanding loudspeakers, are conferred by the electrostatic diaphragm’s large surface area and extremely low mass. The lightweight diaphragm can respond very quickly to input signals, infusing the sound with a lifelike realism and resolution of low-level detail that is generally lacking in dynamic designs. The treble of electrostatic headphones is free from distorting resonances and extends well beyond 20kHz. The downside of electrostatic headphones is the need for a high-voltage outboard power supply to charge the diaphragm, as well as an amplifier of sufficient voltage output to drive the ‘phones. Because this power supply and amplifier must be plugged into an AC wall outlet, electrostatic headphones can’t be used for portable listening.

The classic electrostatic headphones—the model that popularized the technology—are the Stax SR Lambdas. Released in 1979 at the then whopping price of \$340, the SR Lambdas brought unprecedented sound quality to headphones, and was a favorite model among audiophiles and location recording engineers. [Fig.13-6](#) shows a set of contemporary electrostatic headphones.

The third primary headphone technology is the *planar-magnetic* driver: a very thin, lightweight diaphragm to which a conductor is bonded. The diaphragm is suspended in a magnetic field created by permanent magnets on both sides of the diaphragm. The audio signal passes through the conductor bonded to the diaphragm, creating a varying magnetic field. The variations in this magnetic field constitute an analog of the audio signal and interact with the permanent

magnetic field, pushing and pulling the diaphragm to vibrate the air and thus create sound. The drivers in planar-magnetic headphones have much lower moving mass than dynamic drivers, a larger surface area, faster transient response, lower distortion, and greater frequency extension. Planar-magnetics can deliver exceptional sound quality, with resolution, dynamics, and tonal fidelity that rival those of state-of-the-art loudspeakers. Fig.13-7 shows a pair of high-end planar-magnetic headphones.



Fig.13-6 Electrostatic headphones are prized for their lifelike realism and resolution. (Courtesy Stax)



Fig.13-7 Planar-magnetic headphones can deliver state-of-the-art sound quality. (Courtesy Audeze)

Headphone Enhancement Technologies

As great as today's state-of-the-art headphones have become—and they have become spectacularly great—they nonetheless are limited in their ability to create a three-dimensional soundstage. Rather than project instrumental and vocal images in front of the listener within an apparent acoustic space, headphones tend to produce “inside-the-head” imaging. To understand why, consider that when listening to speakers or live sounds, the left ear hears some right-channel information and vice versa, in a phenomenon called interaural crosstalk. The left-channel sound reaching the right ear also arrives slightly later at the right ear than it does at the left ear. But when listening to headphones, the right ear hears *only* right-channel information, and the left ear hears *only* left-channel information. Without interaural crosstalk, the brain is unable to construct a convincing illusion of sounds existing in space in front of us, as we hear with loudspeakers.

Let's look at several techniques and technologies that attempt to make the headphone listening experience more like that of hearing live music:

Crossfeed Circuits

A simple solution to the problem of “inside-the-head” imaging is the *crossfeed circuit*, which has been around since the 1950s. A crossfeed circuit electronically re-creates the interaural crosstalk that occurs when listening to two stereo loudspeakers. The circuit mixes into the left channel some right-channel information that has been greatly attenuated and slightly delayed, and vice versa.

Crossfeed circuits are built into some headphone amplifiers. A switch allows you to turn off the effect if you find the sound unnatural. Many listeners find crossfeed circuits a mixed blessing; they reduce “inside-the-head” imaging, but sometimes at the expense of blurring imaging, softening the treble, reducing resolution, and muting dynamic contrasts. There’s no definitive answer as to whether crossfeed circuits are a benefit; some listeners like the effect, some don’t.

Smyth Virtual Surround (SVS)

The most innovative approach to overcoming the inside-the-head imaging of headphones is undoubtedly Smyth Virtual Surround (SVS). Developed by Stephen Smyth, who invented the codecs that later became the basis for the DTS multichannel surround-sound formats, SVS is a technology that produces through headphones a listening experience indistinguishable from hearing a stereo or multichannel loudspeaker system in a room. SVS is implemented in a commercially available hardware platform called the Smyth Realiser A8.

Here's how it works. You insert in your ear canal earplugs in which are embedded tiny microphones. After connecting the mikes to the Realiser A8 audio processor (a component about the half the width of the average preamplifier), you play a test signal through your home audio system's speakers while you sit in the listening position. You then put on a pair of headphones and play a second test signal, which is again picked up by the mikes inside your ears. The whole process takes about five minutes. Then, whenever you listen to music through headphones connected to the Realiser, you hear a startlingly realistic rendering of *your* speakers in *your* listening room. The imaging is moved outside the headphones into an immersive three-dimensional space that is indistinguishable from the sound of the speakers in the room. During a demo I participated in, I repeatedly removed the headphones and put them back on again, comparing the headphone sound to that of the speakers. For such demonstrations, Smyth has devised a way of instantly and automatically turning on the speaker system when the headphones are removed from the head. I was completely fooled into thinking that the speakers were always playing, even when they weren't.

If you measured different speakers in a different room, the headphone sound will mimic those speakers in that room. As its name suggests, Smyth Virtual Surround works not only for stereo systems, but also for full surround-sound loudspeaker arrays. The impression of hearing surround information located far behind you through stereo headphones is uncanny.

The Realiser A8 incorporates another technology that re-creates through headphones the experience of hearing live music or a loudspeaker system. If you're listening to a live singer in front of you while looking forward, and you turn your head to the right, the image remains in the same position, but now you hear the singer's voice more in your left ear than in your right—and the same is true for listening to a vocalist reproduced by loudspeakers. But when you turn your head while listening through conventional headphones, the image moves

with the rotation of your head. For example, a centrally positioned singer “appears” in whatever direction your head is pointing. This is obviously completely unlike hearing live music, in which aural images are always stationary. Smyth corrects this error by incorporating into the Realiser’s processing algorithm data from a motion-tracking device attached to the headphones. In the demonstration I heard of a multichannel system, I was able to turn in any direction and the soundfield remained fixed, just as I would hear with live music or a loudspeaker system. I did this while taking the headphones on and off several times during each rotation of 360°, and the sound I heard from the headphones was virtually identical to the sound I heard from the multichannel loudspeaker system in every respect, including tonal balance, spatial presentation, and changes in volume.

There’s a reason the room/speaker system and headphone measurements were made with tiny microphones inside your ear canal. To understand that reason, you need to know that the outer ear, called the *pinna*, plays a vital role in locating sounds. The pinna’s pattern of folds and bumps (the medical terms are *folds and bumps*) create a complex series of sound reflections that result in a series of minute delays in the sounds reaching your eardrum. When the direct sounds striking your eardrum are combined with the sounds delayed by the reflections introduced by the pinna, the result is *comb filtering*: a series of notches in the frequency response (comb filtering is described in [Appendix A](#), “Sound and Hearing”). The patterns of the frequency-response notches change with the sound’s direction, and the brain analyzes these comb-filter patterns, and the changes in them, to instantly pinpoint the sound’s direction, if the sound source is moving. The head and torso are also sources of reflected sound and play roles in this mechanism. This modification of a sound by the head and ear is called the *head-related transfer function (HRTF)*.

During its measurement process, the Smyth Realiser A8 captures your particular HRTF (everyone’s is different). On playback, the Realiser uses digital signal processing (DSP) to impose on the audio signal your HRTF, fooling your brain into perceiving the sound as existing in three-dimensional space, complete with the tonal balance and spatial cues present in your room when the measurement was taken. Note that the Realiser also measures and then re-creates the “transfer function” of your loudspeakers and room—the way the speakers and acoustics modify the test signal, and thus the musical signal, on playback.

This technology opens up startling possibilities. You could measure a world-class audio system in a state-of-the-art listening room, then replicate that listening experience through headphones in your own home. Professionals use

the Realiser to replicate, on location or in another studio, a known optimal monitoring environment.

The Realiser A8 is sold as a package with the processing unit, microphones, head-tracking hardware, and a remote control (Fig.13-8). As an option, a pair of Stax electrostatic headphones can be bundled with the Realiser for sales within the U.S. Any headphones can be used—the Realiser will measure and optimize the headphones' response—but the best results will be obtained with high-quality headphones. At the time of this writing, the Realiser A8 package costs \$2910 plus shipping from the U.K., or \$3700 including the Stax SRS-2170 headphones and amplifier.

Had I not experienced SVS firsthand, I wouldn't have thought it possible to produce through headphones the sound of a stereo or multichannel loudspeaker system in a room without the result sounding like a gimmick. Smyth Virtual Surround is an amazing technology that delivers on its promise.



Fig.13-8 The Smyth Realiser A8 can replicate through headphones the sound of loudspeakers in a room. (Courtesy Smyth Research)

Binaural Listening

We've just seen that Smyth Virtual Surround simulates, via digital signal processing, the head-related transfer function that can immerse us in a three-dimensional soundfield when listening to stereo headphones. But there's another way of creating that effect: to record music with an HRTF impressed on the recorded signal. The technique, called *binaural recording*, involves recording with microphones inside the sculpted ears of a dummy head, or *Kunstkopf* (Fig.13-9). When played back through stereo headphones, a binaural recording can re-create the impression of sounds emanating from any location around you. I heard a binaural demonstration of a dummy head getting a "haircut"—in reality, a recording of someone opening and closing a pair of scissors while moving in 360° around the dummy head. The spatial presentation was so realistic that it was unnerving. Note that, to hear the intended spatial effect of binaural recordings, you must listen to them through stereo headphones. Played through loudspeakers, the spatial soundstages of binaural recordings collapse because of interaural crosstalk, the phenomenon of the left ear hearing some sound from the right-channel loudspeaker and vice versa, as described above. Headphones completely isolate each ear from the other ear's channel—a strict requirement for binaural playback.



Fig.13-9 Binaural recordings are made with microphones inserted in the “ear canals” of a dummy head. (Courtesy Georg Neumann)

Binaural recordings can be made without a dummy head: with a pair of microphones placed about 7" apart and facing away from each other, replicating the distance between a human head's ears. These quasi-binaural recordings can provide greater spatial realism through headphones, but don't approach the full spatial realism of which full binaural recordings made with a dummy head are capable.

The ultimate binaural recording would be made using *your* HRTF. The dummy head is just a generic average of all human heads in size and shape, and thus doesn't perfectly deliver precise spatial cues. Everyone's HRTF is different, particularly the comb filtering created by the pinna's folds and bumps. I've read of experiments in which recording microphones were placed inside the ear canals of a human sitting in a concert hall while an orchestra played, and then the recording was later played back to the same person through in-ear monitors. Although playback occurred in a tiny room, the effect was a near-perfect simulation of the spatial experience of sitting in a vast concert hall. As part of this experiment, the recording was played back to listeners other than the one whose HRTF was used to create the recording—and to those listeners, the sound was horribly distorted. This was because the unique HRTFs of the listeners not involved in the recording had been replaced by the HRTF of the single listener who was involved. Each person's brain is adapted to his or her own HRTF; the combination turns what is, objectively, a gross distortion produced by comb filtering into a coherent and natural perception of sound. The uniqueness of one's own adaptation is starkly exposed when one listens to another's head-related transfer function.

Multichannel Headphones

Multichannel playback through stereo headphones can be achieved with less sophisticated technology than Smyth Virtual Surround, though not nearly as convincingly. Often called “Dolby headphones,” this product category uses DSP to mimic the HRTF, enabling the headphones to position sounds anywhere around the listener. Dolby headphones typically include a Dolby Digital decoder that takes a Dolby Digital bitstream from a TV, DVD or Blu-ray Disc player, or other source and outputs a signal to the headphones that creates a facsimile of hearing surround sound through stereo 'phones. With Dolby headphones, you hear dialogue from the center, music and effects at the left and right, and surround information from all around you. Dolby headphones, which are often wireless, are good for watching movies late at night and/or if you don't want to disturb others. They can also be useful for those who are hard of hearing; the greater clarity provided compared to listening to the tiny speakers built into a television can open up a new world. This is particularly true of dialogue, which often gets lost when reproduced by television speakers but is much more intelligible when reproduced through headphones, largely because it's spatially separated from the soundtrack's music and effects.

Headphone Specifications

As with loudspeakers, reading a headphone's specification sheet won't tell you how the headphones sound or how comfortable they are. But, also as with speakers, some headphone specs are useful in selecting the right model for your application.

The most common specification is frequency response, with which we're all familiar. A headphone model's frequency-response spec may read "20Hz–18kHz," but this is misleading. Headphones don't have flat measured response; instead, they're designed to *sound* flat to listeners wearing them. It's not unusual for headphones to have multiple amplitude peaks of as much as 10dB in the upper-midrange to mid-treble region (roughly 3–8kHz). Such peaks in a loudspeaker would make that speaker unlistenable, but in headphones the peaks produce a more natural tonal balance. What's more, frequency-response measurements of headphones are inconsistent compared with speaker measurements; moving the headphone measurement microphone only 1/16" results in a radically different measured response. The measurement protocols for testing freestanding speakers are much more standardized and reliable. Finally, there's no correlation between good speaker frequency response and "good" measured headphone frequency response, because there's no consensus about which headphone frequency response produces the most natural tonal balance. This is why you can't tell much about how a pair of headphones will sound based on looking at their frequency-response spec, or even their frequency-response curve.

A more useful spec is of headphones' impedance. As described in technical detail in [Appendix B](#) ("Audio and Electronics Basics"), impedance is the resistance to current flow (this is a simplified definition). The impedances of headphones span a far wider range of values than do those of freestanding speakers. A loudspeaker's nominal impedance is typically between 4 and 8 ohms, but headphone impedances range from about 15 to 600 ohms.

In [Chapter 5](#) ("Power Amplifiers"), we saw that the lower a speaker's impedance, the more robust the amplifier driving it needs to be. A low-impedance speaker forces the power amplifier to deliver more current to the speaker. That's why a low-impedance speaker is best driven with a powerful amplifier.

Similarly, low-impedance headphones sound best when powered by a high-quality amplifier. The inexpensive amplifiers built into smartphones, tablets, low-end portable music players, and even some A/V receivers can't deliver much output current to the headphones, and are therefore not well suited to driving low-impedance headphones. The sound will be weak and anemic, and the bass will be mushy. The best amplifiers for low-impedance headphones are dedicated outboard headphone amplifiers, or the headphone amplifiers built into high-quality DACs. High-impedance headphones are a better choice with inexpensive mass-market products that include an integral headphone amplifier.

Knowing the headphones' impedance is crucial to matching them with the amplifier that will drive them, specifically the amplifier's output impedance. The rule of thumb is that the headphones' impedance should be at least ten times the amplifier's output impedance. If the amplifier's output impedance is greater than one-tenth the headphones' impedance, the headphones' frequency response (tonal balance) will be affected. This is one reason the same headphones sound different when driven by different amplifiers. In addition, an amplifier with a high output impedance will have a low *damping factor*; that is, its ability to control the motion of the drivers' diaphragms. Low damping factor is associated with lack of tautness, weight, and precision in the bass. For a full technical explanation of this phenomenon, see the sections in [Chapter 5](#) and [Appendix B](#) regarding power-amplifier output impedance. Headphones driven by a headphone amplifier follow the same electrical laws as speakers driven by a home stereo amplifier.

If the headphone amplifier's output impedance is less than 2 ohms, that amplifier should have no problem driving any headphones. But many headphone amplifiers have higher output impedances, especially those found in inexpensive portable devices. The designers of those budget devices add a 30 ohm resistor to the amplifier's output to protect the op-amp driving the headphones. Although this technique makes the amplifier more reliable and stable, it virtually ensures that the device won't sound good with most of the headphones on the market. Many high-end headphone amplifiers have output impedances of a fraction of an ohm, and will drive any headphones with ease. Some high-quality headphones have a very low impedance; their designers assume that they will be driven by low-output-impedance amplifiers.

The last specification you should know about is *sensitivity*: a measure of how much sound the headphone produces for a given amount of input power. Specifically, the sensitivity spec expresses the sound-pressure level (SPL) the headphone outputs when driven by 1 milliwatt (1mW) of power. The sensitivity

of headphones ranges from a low of about 70dB/mW to as high as 125dB/mW. The lower the sensitivity, the more amplifier power needed to achieve a given SPL.

Headphone Amplifiers

Every pair of headphones needs to be driven by an amplifier, whether that amplifier is an integral component in a portable music player or computer, an outboard digital-to-analog converter, the headphone output built into a home-audio component, or a dedicated headphone amplifier. The amplifier's quality plays a large role in how the headphones will sound.

The headphone amplifiers built into most portable devices are inexpensive operational amplifiers (op-amps) that can't generate the voltage and deliver the current required by most headphones. To make matters worse, op-amps' high output impedance introduces many problems, including an erratic frequency response that changes depending on the headphone model, a limited ability to deliver current to the headphones, and low damping factor, all of which contribute to poor bass performance and limited dynamic range.

Higher-quality headphone amplifiers (Fig.13-10) have very low output impedance (as low as 0.01 ohm) and thus a high damping factor, lower distortion, an ability to generate sufficiently high voltage, and are robust enough to deliver current to the headphones. The low output impedance also ensures that the amplifier won't interact with the headphones' impedance curve to introduce frequency-response errors. Virtually all standalone dedicated headphone amplifiers possess these qualities. A few esoteric tubed designs, however, forgo low output impedance in favor of other qualities the designer believes are more important. One example is negative feedback in the amplifier: More negative feedback lowers an amplifier's output impedance, but at the expense of sound quality, in the view of some designers. Such high-output-impedance amplifiers are best used with very-high-impedance (600 ohms) headphones. As always, before purchase, it's best to audition a headphone amplifier with the headphones you plan to use with it.



Fig.13-10 A high-quality dedicated headphone amplifier. (Courtesy Simaudio)

If you choose in-ear monitors, it's important to buy an amplifier with a high signal-noise ratio. The in-ear monitor's high sensitivity converts more of the input signal to sound than do other types of personal listening devices, and that includes any noise generated by the amplifier.

Good amplifier performance isn't strictly the province of dedicated headphone amplifiers; many portable DACs include a high-quality headphone amp. (Sometimes, such products are called "DAC/headphone amplifiers.") Headphone jacks are appearing on an increasing number of products, such as integrated amplifiers and DACs designed for home use. However, the quality of the amplifiers driving those headphone jacks varies greatly with the product. In some components, the headphone jack is merely an afterthought, provided for convenience or to make the product more commercially appealing. In others, the headphone amplifier has received serious design attention and is built to a very high standard. If headphone listening is important to you, search out those products that include a well-designed headphone amplifier.

Balanced and Unbalanced Drive

Most headphone jacks accept the familiar $\frac{1}{4}$ " *phone plug* or the smaller 3.5 or 2.5mm stereo plugs. Each of these plugs has three sections—tip, ring, and sleeve—which is why they're also called *TRS plugs*. The ring and tip respectively carry the right- and left-channel audio signals, and the sleeve is the common ground for both channels. This type of connection, called unbalanced, is shown schematically in Fig.13-11. The right-channel signal, carried by the ring, is connected to one side of the right-channel headphone driver's *voice coil*. The left-channel signal, carried by the tip, is connected to one side of the left-channel headphone driver's voice coil. The sleeve is connected to the other side of both the left and right drivers' voice coils.




Fig.13-11 Unbalanced headphone drive (Courtesy headphone.com)

Some high-end headphones and headphone amplifiers offer balanced connection, sometimes carried via dual three-pin XLR connectors. For portable applications, balanced output is sometimes available on a four-pin XLR connector, or a 3.5 or 2.5mm tip-ring-ring-sleeve (*TRRS*) jack. In balanced connection, the amplifier powers the voice coil in each headphone driver with two audio signals of opposite polarity; the amplifier is connected to both sides of the voice coil. From the discussion of amplifiers in Chapter 5 (“Power and Integrated Amplifiers”) you may recall the technique called *bridging*, in which two amplifier channels drive a loudspeaker, the speaker acting as the “bridge” connecting the two amplifier channels. Balanced headphone drive, shown schematically in Fig.13-12, is identical to bridging in home power amplifiers.

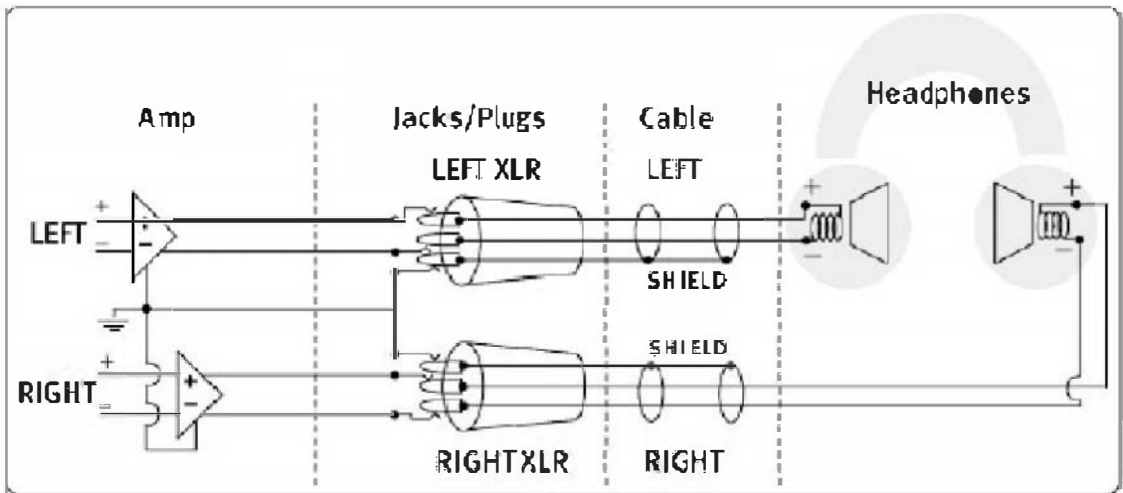


Fig.13-12 Balanced headphone drive (Courtesy headphone.com)

Balanced connection requires that the headphones be wired for balanced drive, and that the amplifier offer a balanced output. The advantages include better control over the diaphragm's movement by the amplifier. This is because a balanced amplifier can swing twice the voltage of an unbalanced amplifier (all else being equal), and swing that voltage in half the time (i.e., a faster slew rate). Distortion is reduced, and *crossstalk* (i.e., unwanted signal leakage between channels) is eliminated because the left and right channels don't share a ground, as they do in unbalanced connection.

Portable Audio

A portable digital-audio player and a pair of headphones is nothing less than a music server and stereo playback system that can fit in your hand. Together, this combination lets you enjoy your favorite music with true high-end sound quality, anywhere and anytime. In my view, that's a compelling reason to invest in a high-quality mobile audio system. This is particularly true with today's portable devices that are compatible with high-resolution audio, that allow you to stream music from lossless subscription services such as Tidal, and that will seamlessly connect with your home music server. A portable audio system can be as simple as a smartphone and a pair of earphones, or as complex as a dedicated player, miniature DAC, outboard headphone amplifier, and high-quality circumaural headphones.

Although it's convenient to store music on a cell phone rather than to also carry a second, music-only device, the phone's sound quality will be limited—the DACs and headphone amplifiers in phones are usually not designed for audiophile-quality sound. A few phones have been designed with high-end audio aspirations, but they don't approach the sound quality possible from a true high-end rig.

There are two alternatives for realizing high-quality music on the go: a laptop computer, or a dedicated personal audio player. The advantages of a laptop are obvious: you're probably carrying the laptop anyway, so there's no need to buy, or haul around, an additional device. But while you can plug headphones directly into a laptop, you'll never get first-rate sound that way. The audio interfaces built into laptops aren't designed for high-quality music reproduction; the DACs are inexpensive, and the headphone amplifier won't adequately drive most headphones.

The solution is a portable DAC and headphone amplifier, such as the models shown in [Fig.13-13](#). The DAC/headphone amplifier plugs into your laptop's USB port, bypassing the laptop's internal DACs and headphone amplifier and substituting its own higher-quality circuitry. The analog output is a 2.5 or 3.5mm stereo jack. Some portable DAC/headphone amps plug directly into your computer's USB port ([Fig.13-13](#), left) with no need for a USB cable. Other portable DACs have a USB port for connecting to the computer via a cable ([Fig.13-13](#), right). Both types are powered from the computer through the USB bus. Virtually all of these products will play files up to 96kHz, but not all will

decode 176.4 and 192kHz files. If you plan to play very-high-resolution files, choose a compatible model. (DACs that support USB 1.0 are limited to 96kHz; DACs that support USB 2.0 can decode files of up to 192kHz.) Portable DACs also vary in the quality of their digital-to-analog conversion circuits and headphone amplifiers. Nonetheless, even the least expensive outboard DAC will sound better than your laptop's integral DAC. In fact, adding this simple and inexpensive device is the single most important step you can take toward better portable sound.



Fig.13-13 A portable DAC and headphone amp can plug directly into your computer's USB port (left), or via a USB cable (right) (Courtesy AudioQuest [left] and Hegel [right])

If you want no-compromise portable sound without the laptop, there's no substitute for a dedicated portable music player. Today's best devices are designed and built like highend DACs and amplifiers, but in miniature. Some incorporate the same DACs found in multi-thousand-dollar home DACs. Their headphone amplifiers are not an afterthought, but an integral and crucial aspect of the design. Today's portable players are fully compatible with high-res files and, when mated with good headphones, offer a surprisingly involving musical experience. When I'm 30,000 feet over the Atlantic, immersed in a private world of my favorite music wonderfully reproduced, I marvel that the experience is even possible.

The iPod may have ushered in the modern era of portable audio, but it faces competition from both quality ends of the market: phones and tablets at the low end, and audiophile-grade portable players at the high end. Let's take a look at what some of today's best portable players offer in technology, features, and sound quality, using as an example the Astell&Kern AK240 (Fig.13-14).

The AK240 is compatible with all PCM formats including DXD, lossless and low-bit-rate codecs (WAV, AIFF, FLAC, ALAC, AAC, OGG, MP3, etc.), as well as standard-rate and double-rate DSD (DSD64 and 128). Built-in WiFi lets you

stream music from the AK240 to your network-connected home-audio system without sonic compromise. This streaming function goes both ways: the AK240 can access and play music stored anywhere on your home network, such as your server's NAS drive. This wireless connectivity also makes possible firmware updates without having to connect the device to a computer or store the update on a microSD card. Dual high-quality DACs and a well-designed headphone amplifier (with the option of balanced outputs) round out the package.



Fig.13-14 Portable players offer high-resolution capability and outstanding sound quality. (Courtesy Astell&Kern)

Things to consider when choosing a portable player are the size of the player's internal memory for music storage, and if that memory capacity can be expanded (usually by inserting a microSD card). Keep in mind that high-res WAV files can quickly consume even high-capacity internal memory. It seems as though there's a universal law of music servers that says you'll always, eventually, run out of storage.

If you want to take the next step in building a personal portable music system, consider adding an external headphone amplifier to the portable music player. The better external headphone amps sound better than a portable player's

integral amp. The downsides are the additional expense and complexity, and having another device to lug around. Audition the portable player and your headphones together, then try a separate amplifier to gauge the potential improvement.

Desktop Audio

The computer-audio revolution has provided a simple path for anyone to enjoy music on the desktop while working. You likely already use a computer running iTunes that is connected to the Internet's streaming music services, so why not add a small amplifier and a pair of desktop speakers? In fact, I'm listening to my desktop-audio system as I write this.

The simplest way to get desktop audio is by adding a pair of small speakers with integral DACs and amplifiers. Connect the speakers to your computer via a USB cable, optical connection (TosLink), or S/PDIF coaxial (on an RCA jack), and you're finished.

Note that powered "computer speakers" don't contain an integral DAC; they're fed from the computer's analog output. As we saw in the section on portable audio, you'll never get good sound from the computer's analog output, which is why computer speakers can be greatly improved by adding an external DAC.

The next step up in quality is realized with passive desktop loudspeakers, an outboard DAC, and amplifier. The portable DAC/headphone amplifiers described in the previous section will also work in a desktop system. But rather than driving headphones, the DAC's analog output connects to a small amplifier that will drive the desktop speakers. Because nearly all desktop amplifiers have RCA input jacks, you'll need a 3.5mm-to-RCA breakout cable (Fig.13-15).



Fig.13-15 A breakout cable with a 3.5mm stereo plug at one end and dual RCA jacks at the other end is often needed between a DAC and desktop amplifier. (Courtesy AudioQuest)

DACs designed for desktop audio have RCA outputs, obviating the need for the breakout cable. They also differ from their portable brethren in that they're powered from an AC wall outlet (often through a "wallwart" transformer) rather than from the computer via the USB bus. Powering a DAC from an AC wall outlet rather than the USB bus is a big sonic advantage because the power to the DAC circuitry can be cleaner and more stable. Desktop DACs may also include a volume-control knob.

High-quality DACs can also be built into integrated amplifiers, replacing two components on your desktop with one (Fig.13-16). Although the integrated amplifiers with built-in DACs designed for driving floorstanding speakers in a living room will work on a desktop, they tend to be big and expensive, and are often overkill for desktop audio. A number of companies offer small, affordable, unobtrusive integrated amps with high-quality DACs and power outputs that range from about 20 to 100Wpc. A few provide an output jack for connecting a powered subwoofer. Many are built around Class D output stages, which offer small size, low cost, light weight, and cool operation.

Desktop speakers are typically small, two-way designs that have been specifically designed for *nearfield* listening (Fig.13-17). Smaller bookshelf speakers will also work on a desktop. Some desktop-specific speakers are angled so that the tweeters are aimed up toward your ears. Like any other speaker, passive desktop speakers need to be matched to the amplifier driving them. Small speakers tend to have low sensitivity, which requires more amplifier power. But because you'll usually sit just a few feet away from the speakers, they don't need to play very loudly.

If you want more bass than what's possible from a speaker that fits on a desktop, you may choose to add a subwoofer. The guidelines in Chapter 16 ("System Setup Secrets") for selecting and setting up a subwoofer for your living room also apply to subwoofers that augment the bass output of desktop speakers. It's crucial that the signal driving the desktop speakers has been high-pass-filtered by a crossover, either in the integrated amplifier driving them or in the subwoofer itself.

Between your home system, portable audio player and headphones, and desktop audio package, you'll have great-sounding music whenever and wherever you want.



Fig. 13-16 A combination DAC/amplifier designed for desktop audio (Courtesy CEntrance)



Fig.13-17 Some desktop loudspeakers can offer outstanding performance. (Courtesy Audience)

14 System Setup Secrets, Part One: Loudspeaker Placement and Room Acoustics

Introduction

The room in which music is reproduced, and the positions of the loudspeakers within that room, have a profound effect on sound quality. In fact, the loudspeaker/listening room interface should be considered another component in the playback chain. Because every listening room imposes its own sonic signature on the reproduced sound, your system can sound its best only when given a good acoustical environment. An excellent room can help get the most out of a modest system, but a poor room can make even a great system sound mediocre.

Fortunately, you can greatly improve a listening room with a few simple tricks and devices. The possibilities range from simply moving your loudspeakers—or even just your listening seat—a few inches, to building a dedicated listening room from scratch. Between these two extremes are many options, including adding inexpensive and attractive acoustical treatment products.

This chapter covers optimizing your system acoustically, beginning with the most effective and least expensive acoustical technique of all—speaker placement. Keep in mind that this chapter describes how to set up a 2-channel stereo, not a home theater or multichannel system. Advice on setting up those systems is in [Chapters 11 and 12](#). Don't be intimidated by the technical information toward the end of the chapter—it's included for those wanting to know some of the theory behind the practical information presented in the beginning sections. (If you find

this later section difficult, I've summarized it under "Acoustical Dos and Don'ts.")

Loudspeaker Placement

The most basic problem in many listening rooms is poor loudspeaker placement. Finding the right spot for your speakers is the single most important factor in getting good sound in your room. Speaker placement affects tonal balance, the quantity and quality of bass, soundstage width and depth, midrange clarity, articulation, and imaging. As you make large changes in speaker placements, then fine-tune these positions with smaller and smaller adjustments, you'll hear in the sound a newfound musical rightness and seamless harmonic integration. When you get it right, your system will come alive. Best of all, it costs no more than a few hours of your time.

Before getting to specific recommendations, let's cover the six fundamental factors that affect how a loudspeaker's sound can change with their positions. Later we'll look at each of these factors in detail.

1) *The relationship between the loudspeakers and the listener is of paramount importance.* The listener and speakers should form an equilateral triangle; without this basic setup, you'll never hear good soundstaging and imaging.

2) *Proximity of loudspeakers to walls affects the amount of bass.* The nearer the loudspeakers are to walls and corners, the louder the bass.

3) *The loudspeaker and listener positions in the room affect the audibility of room resonant modes.* *Room resonant modes* are reinforcements of certain frequencies that create peaks and dips in the frequency response, which can add an unnatural "boominess" to the sound. When room resonant modes are less audible, the bass is smoother, better defined, and midrange clarity increases. (Room resonant modes are described in detail later in this chapter.)

4) *The farther out into the room the loudspeakers are, the better the soundstaging—particularly depth.* Positioning loudspeakers close to the wall behind them can destroy the impression of a deep soundstage.

5) *Listening height affects tonal balance.* With some loudspeakers, how high your ears are in relation to the speakers' tweeters can affect the amount of treble you hear.

6) Toe-in (angling the loudspeakers toward the listener) affects tonal balance (particularly the amount of treble), soundstage width, and image focus. Toe-in is a powerful tool for dialing in the soundstage and treble balance.

Let's look at each of these factors in detail.

1) Relationship between the loudspeakers and the listener

The most important factor in getting good sound is the geometric relationship between the two loudspeakers and the listener (we aren't concerned about the room yet). The listener should sit exactly between the two speakers, at a distance away from each speaker that's slightly greater than the distance between the speakers themselves. Though this last point is not a hard-and-fast rule, you should certainly sit exactly between the loudspeakers; that is, the same distance from each one. If you don't have this fundamental relationship, you'll never hear good soundstaging from your system.

[Fig.14-1](#) shows how your loudspeaker and listening positions should be arranged. The listening position—equidistant from the speakers, and slightly farther from each speaker than the speakers are from each other—is called the “sweet spot.” This is roughly the listening position where the music will snap into focus and sound the best. If you sit to the side of the sweet spot, the soundstage will tend to bunch up around one speaker. This bunching-up effect will vary with the speaker model; some pairs of speakers produce wider sweet spots than others.




Fig.14-1 The listener should sit precisely between the loudspeakers, and at the same distance from each speaker.

Setting the distance between the speakers is a trade-off between a wide soundstage and a strong center, or phantom, image. The farther apart the loudspeakers (assuming the same listening position), the wider the soundstage. As the speakers are moved farther apart, however, the center image weakens, and can even disappear. If the speakers are too close together, the soundstage narrows.

Speakers placed the optimal distance apart will produce a strong center image and a wide soundstage. You can experiment with the effective speaker width simply by moving your listening chair forward and back. There will likely be a position where the center image snaps into focus, appearing as a stable and almost tangible presence exactly between the speakers. A musical selection with a singer and sparse accompaniment is ideal for setting loudspeaker spacing and ensuring a strong center image. With the speakers fairly close together, listen for a tightly focused image exactly between the two speakers. Move the speakers a little farther apart and listen again. Repeat this move/listen procedure until you start to hear the central image become larger, more diffuse, and less focused, indicating that you've gone slightly beyond the maximum distance your speakers should be

from each other for a given listening position. Wait until after you've completed the entire speaker-placement procedure to install the speakers' spikes.

A factor to consider in setting this angle is the relationship of this triangle to the room. You can have the same geometric relationship between speakers and sweet spot with the speakers close together and a close listening position, or with the speakers farther apart and a more distant listening position. At the farther listening position, the listening room's acoustic character will affect the sound more than at the nearer position. That's because you hear more direct sound from the speakers and less reflected sound from the room's walls. Consequently, the farther away you sit, the more spacious the sound. The closer you sit, the more direct and immediate the sound. A technique that's helpful in small rooms, or those with fundamentally poor acoustics, is to sit very close to the speakers, a condition called *near field* listening. Some speakers need a significant distance between them and the listener to allow the outputs of the speakers' individual drive units to integrate. If you hear a large tonal difference just by sitting closer, you should listen from a point farther away from the speakers.

2) Proximity to walls affects the amount of bass

Room boundaries have a great effect on a speaker's overall tonal balance. Loudspeakers placed close to walls will exhibit a reinforcement in the bass (called "room gain"), making the sound more bottom-heavy. Some speakers are designed to be placed near the room's front wall (the wall behind the speakers); they need this reinforcement for a more natural tonal balance. Such speakers sound thin if placed out into the room. Others sound thick and heavy if not at least 3' from the front and sidewalls. Be sure which type you're buying if your placement options are limited.

When a loudspeaker is placed near a wall, more of its bass energy is reflected back into the room, essentially in phase with the speaker's direct output. This means that the direct and reflected waves reinforce each other at low frequencies, producing louder bass. Fig.14-2 shows the difference in a speaker's frequency response when measured in an anechoic chamber (a room that reflects no soundwaves) and in a normal room. A frequency-response graph plots amplitude (loudness) vs. frequency. As you can see in the graph, not only is the bass boosted by the room, the speaker's low-frequency extension is deepened. Each surface near the speaker (floor, front wall, and sidewalls) will increase the speaker's bass output. The closer to the room's corners the speakers are, the more bass you'll hear.




Fig.14-2 Comparison of a loudspeaker's anechoic and in-room responses

The speaker's position in relation to the front and sidewalls will also affect which frequencies are boosted. Correct placement can not only extend a loudspeaker's bass response by complementing its natural rolloff, but can prevent peaks and dips in the response. Improper placement can cause frequency-response irregularities that color the bass. That is, some frequencies are boosted relative to others, making the bass reproduction less accurate. The graph of Fig.14-3a is a speaker's in-room response (the speaker's response as modified by the listening room) when placed equidistant from the front and sidewalls. Note the 10dB notch (reduction in energy) at about 200Hz and the peak centered at 60Hz. The result will be boominess in the bass and leanness in the midbass. Moving the speaker different distances from the front and sidewalls can make the response much smoother (Fig.14-3b). The same speaker's response at a different room position is shown in Fig.14-4.




Fig.14-3a Placing the speakers the same distance from the front and sidewalls produces peaks and dips in the bass response.




Fig.14-3b Placing the speakers different distances from the front and sidewalls smooths the bass response.




Fig.14-4 Response of the same speaker in Figs. 14-3a and 14-3b, but in a different room position (Figs. 14-2, 14-3a, 14-3b, and 14-4 Courtesy Avalon Acoustics, Inc.)

These graphs illustrate that a loudspeaker’s position in the room affects the bass response, and that the loudspeakers should be positioned at different distances from the front and sidewalls. A rule of thumb: the two distances shouldn’t be within 33% of each other. For example, if the speaker is 3’ from the sidewall, it should be at least 4’ from the front wall.

Many speaker manufacturers will specify a distance from the front and sidewalls. When a measurement is specified, the distance is between the woofer cone and the wall. Start with the speakers in the manufacturer-recommended positions, then begin experimenting.

How close the speakers are to the sidewalls affects the amplitude of the sidewall reflection. The closer the speakers are to the sidewalls, the higher the level of the sidewall reflections reaching the listener—not a good thing. If you’ve treated the sidewalls as described later in this chapter, putting the speakers closer to the sidewalls won’t have as great an effect as if the sidewalls were left untreated.

A highly effective technique for finding the best distances from the front and sidewalls was developed by David Wilson, of Wilson Audio Specialties. I’ve watched him use this technique a number of times. Stand at the front wall, about the same distance from the sidewalls as you expect the speakers to end up. Begin speaking, moving in a line perpendicular to the front wall. Listen to the timbre of your voice as you move. You’ll hear different colorations as you move, and then, suddenly, at one particular location, your voice will sound open, clear, and

uncolored. There will be a small range of several inches over which your voice sounds the most natural. Mark this area on the floor with masking tape. This is the best distance from the front wall for your speakers. Now repeat this exercise, but starting on the sidewall at the same plane at which you have placed the masking tape. The point at which your voice sounds clearest is the best distance from the sidewall to place your speaker. The two points will intersect, indicating where you should place the speaker. The intersection should be just behind the loudspeaker's front baffle and centered on the woofer. This technique is highly effective, and correlates very well with the placements suggested by computer modeling. The accuracy can be improved by having someone sit in the listening position as you move and speak, and confirm the point at which your voice sounds least colored.

3) Loudspeaker and listener positions affect room-mode audibility.

In addition to deepening bass extension and smoothing bass response, correct speaker placement in relation to the room's walls can also reduce the audible effects of your room's *resonant modes*. Room resonant modes, described later in this chapter, are reinforcements at certain frequencies that create peaks in the frequency response. Room modes create *standing waves*, which are stationary patterns of high and low sound pressure in the room that color the sound. The standing-wave patterns in a room are determined both by the room's dimensions and by the position of the sound source in the room. By putting the speakers and listener in the best locations, we can achieve smoother bass response.

A well-known rule of thumb states that, for the best bass response, the distance between the speakers and the front wall should be one-third the length of the room (Fig.14-5). If this is impractical, try one-fifth the room length. Both of these positions reduce the excitation of standing waves and help the sounds from the speakers integrate with the room. Ideally, the listening position should be two-thirds of the way into the room.

Starting with these basic configurations, move the loudspeakers and the listening chair in small increments while playing music rich in low frequencies. Listen for smoothness, extension, and how well the bass integrates with the rest of the sound. When you find the speaker position where the bass is smoothest, you should also hear an increase in midrange clarity and definition.




Fig.14-5 The "rule of thirds" says to place the speakers and listening position a third of the way into the room. If this isn't practical, try one-fifth of the way into the room.

An excellent test signal for evaluating bass and midrange sound in a room is the Music Articulation Test Tone (MATT), developed by Acoustic Sciences Corporation (ASC). This special test signal is a series of tone bursts that rise in pitch, with silences between the bursts. Ideally, you should hear the bursts and silences as separate events. When heard through headphones or with your ear near the speaker, each burst is clearly articulated. But when the sound is modified by the listening room, certain frequency bands of the ascending tone bursts become smeared or garbled, indicating that the listening room is storing, and then releasing, energy at those frequencies. By moving the speakers and listening to the MATT, you can easily discover where your speakers work best in your room. (The MATT is available on *Stereophile's Test CD 2*, which also includes more detailed information about how to use this unique test signal. This disc is available at tubetrap.com.)

Another way of finding the spots for your loudspeakers and listening chair that best reduce the influence of standing waves is a computer program called Room Optimizer, created by RPG Diffusor Systems. The program asks for your room dimensions, loudspeaker type and woofer location within the cabinet, and

other information. The software performs multiple calculations to determine the optimal speaker and listener positions. Room Optimizer is highly recommended for this aspect of speaker placement.

4) Distance from front wall affects soundstaging

Generally, the farther away from the front wall the speakers are, the deeper the soundstage. A deep, expansive soundstage is rarely developed with the loudspeakers near the front wall; pulling the speakers out a few feet can make the difference between poor and spectacular soundstaging. Unfortunately, many living rooms can't accommodate speakers far out into the room. If the speakers must be close to the front wall, make that wall acoustically absorbent.

5) Listening height and tonal balance

Most loudspeakers exhibit changes in frequency response with changes in listening height. These changes affect the midrange and treble, not the bass balance. Typically, the sound will be brightest (i.e., have the most treble) when your ears are at the same height as the tweeters, or on the tweeter axis. Most tweeters are positioned between 32" and 40" from the floor to coincide with the height of the ears of the typical seated listener. If you've got an adjustable office chair, you can easily hear the effects of listening axis on tonal balance.

The degree to which the sound changes with height varies greatly with the loudspeaker. Some models have a broad range over which little change is audible; others can exhibit large tonal changes when you merely straighten your back while listening. Choosing a listening chair that puts your ears on the optimal axis will help achieve a good treble balance.

This difference in response is easily measurable. A typical set of loudspeaker measurements will include a family of response curves measured on various axes. The on-axis response (usually on the tweeter axis) is normalized to produce a straight line; the other curves show the *difference* between the on-axis response at various heights. [Chapter 16](#) ("Specifications and Measurements") includes a full description of how to interpret these measurements.

6) Toe-in

Toe-in is pointing a loudspeaker inward toward the listener rather than aiming it straight ahead (see [Fig.14-6](#)). There are no rules for toe-in; the optimal amount will vary greatly with the speaker and the room. Some speakers need toe-in;

others work best firing straight ahead. Toe-in affects many aspects of the sound, including mid- and high-frequency balance, soundstage focus, sense of spaciousness, and immediacy.




Fig.14-6 Loudspeakers positioned with no toe-in (left) and with toe-in (right)

Most loudspeakers sound brightest when listened to directly on-axis (directly in front of the loudspeaker). Toe-in thus increases the amount of treble heard at the listening seat. An overly bright speaker can often be tamed by pointing it straight ahead. Some models, designed for listening without toe-in, are far too bright when listened to on-axis.

The ratio of direct to reflected sound increases with toe-in. That's because a toed-in loudspeaker will present more direct energy to the listener and project less energy toward the room boundaries, where it might reach the listener only after being reflected from those surfaces. In a listening room with reflective sidewalls, toeing-in the speakers can be a decided advantage. Moreover, the amplitude of sidewall reflections is greatly decreased with toe-in. Conversely, less toe-in increases the amount of reflected energy heard by the listener, adding to a sense of spaciousness and air. Reducing toe-in can open up the soundstage and create a feeling of envelopment.

Similarly, toe-in often increases soundstage focus and image specificity. When toed-in, many loudspeakers provide a more focused and sharply delineated soundstage. Images are more clearly defined, compact, and tight, rather than

diffuse and lacking a specific spatial position. The optimal toe-in angle is often a trade-off between too much treble and a strong central image. With lots of toe-in, the soundstage snaps into focus, but the sound is often too bright. With no toe-in, the treble balance is smoother, but the imaging is more vague.

Toe-in also affects the sound's overall spaciousness. No toe-in produces a larger, more billowy, less precise soundstage. Instruments are less clearly delineated, but the sound is bigger and more spacious. Toeing-in the speakers shrinks the apparent size of the soundstage, but allows more precise image delineation. Again, the proper amount of toe-in depends on the loudspeaker, the room, and personal preference. There's no substitute for listening, adjusting toe-in, and listening again.

Identical toe-in for both speakers is essential to soundstaging. A speaker's frequency response changes with toe-in, and identical frequency response from both speakers is crucial to the precise placement of images in the soundstage. Achieving identical toe-in can be accomplished by measuring the distance from the front wall to both rear corners of one speaker; these distances will differ according to the degree of toe-in. Repeat this procedure on the other speaker, adjusting its toe-in so that the distances match those of the first speaker. Another way to ensure identical toe-in is to sit in the listening seat and look at the speakers' inside edges. You should see the same amount of each speaker cabinet's inner side panel. You can also use a laser-alignment tool to ensure identical toe-in. Substituting a piece of cardboard or similar material for the listener in the listening seat, mount the laser on the top of one of the speakers, flush with one cabinet edge. Mark on the cardboard where the beam hits. Repeat the process with the other speaker and adjust the toe-in so that the beam strikes the same spot as the first beam. The laser alignment tool will come in handy later to verify that each speaker has the same degree of backward tilt (or no tilt at all). The degree of tilt, called the rake angle, can vary if the speaker's spikes are not screwed into the speaker at uniform depths.

Keep in mind that any change in one parameter of speaker placement will affect all other parameters. For example, a wide soundstage can be achieved with narrow placement but no toe-in, or wide placement with extreme toe-in.

Loudspeaker Placement in Asymmetrical Rooms

So far I've discussed positioning speakers in an ideal room that is symmetrical and bounded on four sides. But what about those listening rooms in which one side is open to the rest of the house, or rooms in which proper speaker placement is difficult? Fortunately, you can still get good sound in an odd-shaped room.

Let's first consider the room in which one wall is essentially missing, where the listening room opens on the right to the rest of the house (Fig.14-7). The left speaker will introduce more room gain than the right speaker because of its proximity to the corner. Remember that the closer a speaker is to walls, the more bass you'll hear. The right speaker has much less bass reinforcement, creating a bass imbalance.

The solution is twofold. First, keep the speakers fairly far from the front wall. This minimizes the room gain for the left loudspeaker, making its response more closely match that of the right speaker. Second, install a bass absorber (such as an ASC Tube Trap) in the corner behind the left speaker. This trap will absorb bass energy rather than reflect it back into the room, and thus will mimic the lack of room reinforcement from the right speaker. You'll never be able to get a precise match, but these two techniques should get the two speakers close in bass output. I had a listening room for about two years that looked like Fig.14-7, and got good results with a 16" ASC Full Round Tube Trap in the left corner, behind the speaker.




Fig.14-7 A room in which one side is open can cause a bass imbalance.

If a room has nonparallel walls, try to position the speakers at the room's narrow end. Figs.14-8a and b show incorrect and correct speaker placement in such a room. The problem with the placement in Fig.14-8a (left) is that detrimental sidewall reflections are directed at the listening position. In Fig.14-8b (right), sidewall reflections are naturally directed away from the listener. Many recording-studio control rooms are narrower at the speaker end for just this reason.

Short-Wall vs. Long-Wall Placement

The examples shown have all assumed that the loudspeakers will be placed along one of the room's short walls. This traditional method provides the most space for bringing the speakers out into the room and away from the front wall. But putting the speakers on a long wall (Fig.14-9) also has some advantages.

Long-wall placement results in more direct sound at the listening position and less sidewall-reflected energy. This confers a large advantage in tonal purity and soundstage accuracy, for reasons described later in this chapter. You hear more from the speakers and less from the room. The sound tends to be more immediate, detailed, and present.




Fig.14-8 In angled rooms, speakers should be placed so that sidewall reflections are directed away from the listener.

Long-wall placement works only when the room is wide enough to accommodate some space between the speakers and the wall behind them, space

between the speakers and listening position, and some space between the listener and room's rear wall. If you sit with your head very near the wall behind you, the sound will likely be boomy. Consequently, you'll end up sitting fairly close to the speakers with long-wall placement in all but very large rooms. As described in [Chapter 6](#), speakers with first-order crossovers require significant space between speakers and listener for the sounds from the individual drivers to properly integrate. If you sit too close to a speaker with first-order crossovers, the tonal balance will be wrong. Consider the room's width and the type of crossovers in your loudspeakers when deciding if long-wall placement is right for you. The reduction in reflected sound at the listening position with long-wall placement is a considerable advantage




Fig.14-9 Positioning the loudspeakers along the room's long wall reduces sidewall reflections at the listening position.

Finally, you may consider placing your loudspeakers at an odd angle in relation to the room, as shown in [Fig.14-10](#). I've seen speaker designers use this technique at hi-fi shows when trying to get good sound in a small hotel room. This placement, which works best with small speakers, has the advantage of reducing sidewall reflections at the listening position. The disadvantages are that it consumes much of the room's floor space and positions the speakers closer to walls, where the bass will be reinforced. That's why this technique is more successful with small speakers that have limited bass extension.

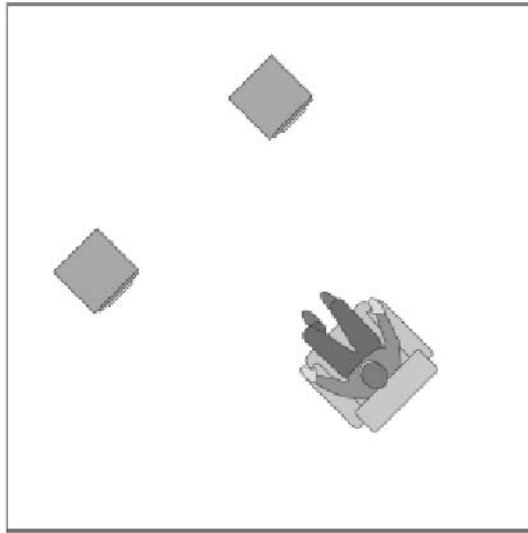


Fig.14-10 Angled loudspeaker placement works best with small speakers.

Dipolar and Bipolar Loudspeaker Placement

Dipolar loudspeakers produce sound to the rear as well as to the front. An electrostatic speaker is a dipole because its vibrating diaphragm sits in open space rather than in a cabinet, launching sound equally to front and rear. This rear wave from dipolar speakers is *out of phase* with the front wave; that is, when the diaphragm moves forward to create positive pressure in front of the diaphragm, it creates negative pressure behind the diaphragm.

A *bipolar* speaker typically uses arrays of conventional dynamic drivers on the front and rear of the loudspeaker enclosure. The front and rear waves from a bipolar speaker are *in phase* with each other. That's the difference between dipolar and bipolar: a dipole's rear wave is *out of phase* with the front wave, and a bipole's rear wave is *in phase* with the front wave.

Both of these speaker types are covered in detail in [Chapter 6](#). What concerns us here are the special placement requirements of bipolar and dipolar speakers, along with the different ways in which they interact with the listening room.

The most important consideration when positioning dipoles is that the wall behind the speakers (the front wall) has a much greater influence on the sound than it does with conventional point-source speakers (those that direct energy in only one direction). Conversely, how the sidewalls are treated is less important with dipoles because they radiate very little energy to the sides. (See [Fig.6-5](#) in [Chapter 6](#), which shows the dispersion patterns of point-source and dipolar speakers.)

Generally, dipoles like a reflective front wall, but with some diffusing objects behind the speaker to break up the reflected energy. A highly absorbent front wall defeats the purpose of a dipole; that reflected energy is beneficial, and you want to hear it. But if the wall is flat and lacks surfaces that scatter sound, the reflected energy combines with the direct sound in a way that reduces soundstage depth. Bookcases directly behind dipolar speakers help diffuse (scatter) the rear wave, as do rock fireplaces, furniture, and other objects of irregular shape. ASC Tube Traps can be positioned directly behind dipoles with their reflective sides out, or you can go all the way with RPG Diffusors. RPG's Skyline diffuser ([Fig.14-11](#)) is ideal for diffusing the rear wave from dipoles. These 2' by 2' squares are highly effective in scattering sound, and can be attached to a wall or mounted on a stand. Tall, dipolar speakers benefit from two or three Skylines arrayed vertically directly behind each speaker.




Fig.14-11 RPG's Skyline diffuser effectively scatters sound. (Courtesy RPG Diffusor Systems, Inc.)

You still want to treat the sidewalls to absorb reflections with dipoles, but their narrow radiation pattern makes treating the sidewalls less important than with point-source speakers. This is particularly true of dipoles that sound their best toed-in toward the listening position. Similarly, dipoles have very narrow vertical dispersion, meaning they direct very little sound toward the ceiling and floor.

Dipolar loudspeakers also need to be placed farther out into the room than conventional point-source speakers. You can't put dipoles near the front wall and expect a big, deep soundstage. Be prepared to give up a significant area of your listening room to dipolar speakers.

Subwoofer Setup and Placement

It's relatively easy to put a subwoofer into your system and hear more bass. What's difficult is making the subwoofer's bass *integrate* with the sound of your main speakers. Low bass as reproduced by a subwoofer's big cone can sound different from the bass reproduced by the smaller cones in the left and right speakers. A well-integrated subwoofer produces a seamless sound, no boomy thump, and natural reproduction of music. A poorly integrated subwoofer will sound thick, heavy, boomy, and unnatural, calling attention to the fact that you have smaller speakers reproducing the frequency spectrum from the midrange up, and a big subwoofer putting out low bass.

Integrating a subwoofer into your system is challenging because the main speakers may have small cones, and the subwoofer has a large, heavy cone. Moreover, the subwoofer is optimized for putting out lots of low bass, not for reproducing detail. The main speakers' upper bass is quick, clean, and articulate. The subwoofer's bass is often slow and heavy. The word *slow* is not technically correct, but vividly describes how such a subwoofer sounds. More precisely, a "slow" subwoofer is one that has excessive *overhang*: the cone keeps moving long after the drive signal has stopped, which conveys the impression that the bass is "slow."

Achieving good integration between small speakers and a subwoofer is easier if you buy a complete system made by one manufacturer. Such systems are engineered to work together to provide a smooth transition between the subwoofer and the main speakers. Specifically, the crossover network removes bass from the left and right speakers, and removes midrange and treble frequencies from the signal driving the subwoofer. If all these details are handled by the same designer, you're much more likely to get a smooth transition than if the subwoofer is an add-on component from a different manufacturer.

If you do choose a subwoofer made by a different manufacturer, several controls found on most subwoofers help you integrate the sub into your system. One control lets you adjust the *crossover frequency*. This sets the frequency at which the transition between the subwoofer and the main speakers takes place. Frequencies below the crossover frequency are reproduced by the subwoofer; frequencies above the crossover frequency are reproduced by the main speakers. If you have small speakers that don't go very low in the bass and you set the crossover frequency too low, you'll get a "hole" in the frequency response. That is,

there will be a narrow band of frequencies that aren't reproduced by the woofer *or* the main speakers.

Setting the subwoofer's crossover frequency too high also results in poor integration, but for a different reason. The big cone of a subwoofer is specially designed to reproduce low bass. When it is asked to also reproduce upper-bass frequencies, those upper-bass frequencies are less clear and distinct than if they were reproduced by the smaller main speakers. Finding just the right crossover frequency is the first step in achieving good integration. Most subwoofer owner's manuals include instructions for setting the crossover frequency. As a rule of thumb, the lower the subwoofer's crossover is set, the better.

Some subwoofers also provide a knob or switch marked Phase. Phase, described in detail in [Appendix A](#), is a time difference between two soundwaves. To understand a subwoofer's phase control, visualize a sound wave being launched from your subwoofer and from your main speakers at the same time. Unless the main speakers and subwoofer are identical distances from your ears, those two soundwaves will arrive at your ears at different times; that is, they will have a phase shift between them. In addition, the electronics inside a subwoofer can create a phase shift in the signal. The Phase control lets you delay the wave generated by the subwoofer so that the subwoofer's wave is perfectly in-sync with the wave from the main speaker. When the soundwaves are in phase, you hear a more coherent and better-integrated sound. One way of setting the phase control is to sit in the listening position with music playing through the system. Have a friend rotate the phase control (or flip the phase switch) until the bass sounds the smoothest.

But there's a much more precise way of setting the phase control that guarantees perfect phase alignment between the subwoofer and main speakers. First, reverse the connections on your main loudspeakers so that the black speaker wire goes to the speaker's red terminal, and the red speaker wire goes to the speaker's black terminal. Do this with both speakers. Now, from a test CD that includes pure test tones, select the track whose frequency is the same as the subwoofer's crossover frequency. Sit in the listening position and have a friend rotate the subwoofer's phase control until you hear the *least* amount of bass. The subwoofer's phase control is now set perfectly. Return your speaker connections to their previous (correct) positions: red to red, black to black.

Here's what's happening when you follow this procedure. By reversing the polarity of the main speakers, you're putting them out of phase with the subwoofer. When you play a test signal whose frequency is the same as the subwoofer's crossover point, both the sub and the main speakers will be

reproducing that frequency. You'll hear minimum bass when the waves from the main speakers and subwoofers are maximally out of phase. That is, when the main speakers' cones are moving in, the subwoofer's cone is moving out. The two out-of-phase waves cancel each other, producing very little bass. Now, when you return your main speakers to their proper connection (putting them back in phase with the subwoofer), they will be maximally *in-phase* with the subwoofer. This is the most accurate way to set a subwoofer's phase control, because it's easier to hear and identify the point of maximum cancellation than the point of maximum reinforcement. Unless you later move the subwoofer or main speakers, you need to perform this exercise only once.

The best integration results from adding two (or more) subwoofers to your system. Two subwoofers drive the air in the room more uniformly, with fewer extremes of standing waves. The result is smoother bass throughout the room, and better integration with the main speakers.

You can also get more dynamic impact and clarity from your subwoofer by placing it close to the listening position. Sitting near the subwoofer causes you to hear more of the sub's direct sound and less of the sound after it has been reflected around the room. You hear—and feel—more of the low-frequency wave launch, which adds to the visceral impact.

Subwoofer placement also has a large effect on how much bass you hear and how well the sub integrates with your main speakers. When a subwoofer is correctly positioned, the bass will be clean, tight, quick, and punchy. A well-located subwoofer will also produce a seamless sound between the sub and the front speakers; you won't hear the subwoofer as a separate source of sound. A poorly positioned subwoofer will sound boomy, excessively heavy, thick, lacking detail, slow, and have little dynamic impact. In addition, you'll hear exactly where the front speakers leave off and the subwoofer takes over.

The simplest, most effective way of positioning a subwoofer is to temporarily put it as close to the listening position as feasible. Raise the subwoofer off the floor, if possible, so that it's close to where the listeners' ears will normally be. Play a piece of music with an ascending and descending bass line, such as a "walking" bass in straight-ahead jazz. Crawl around the floor on your hands and knees (make sure the neighbors aren't watching) until you find the spot where the bass sounds smoothest, and where each bass note has about the same volume and clarity. Avoid positions where some notes last longer, and/or sound slower or thicker, than others. When you've determined where the bass sounds best, put the subwoofer there permanently. Now, when you're back in the listening seat, the bass should sound smooth and natural.

Some general guidelines for subwoofer placement: As with full-range speakers, avoid putting the subwoofer the same distance from two walls. For example, if you have a 20'-wide room, don't put the subwoofer 10' from each wall. Similarly, don't put the subwoofer near a corner and equidistant from the side and front walls. Instead, stagger the distances to the walls. Staggering the subwoofer's distance from each wall smooths the bass because the frequencies being reinforced by the wall are randomized rather than coincident.

Multichannel Loudspeaker Placement

So far, we've discussed the placement of two loudspeakers for stereo music reproduction. With multichannel music and home theater becoming increasingly common, let's expand on these speaker-placement principles to include positioning more than two speakers.

As described in [Chapter 12](#) ("Multichannel Audio"), the optimal loudspeaker types (point source, dipolar) and configurations for multichannel music reproduction and for home theater are different. For multichannel music, the ideal speaker array is five identical full-range speakers placed equidistant from the listener. For film-soundtrack reproduction, the center speaker is typically smaller and lacks bass extension, and the surround speakers are dipolar types mounted on the sidewalls. Speaker arrays optimized for home theater also include a subwoofer.

Let's first look at the ideal loudspeaker array for multichannel music. The center speaker should be positioned on the room's center-line directly in front of the listening position, and slightly behind the plane described by the forwardmost edges of the left and right loudspeakers. This placement creates a gentle arc, and puts the center speaker at the same distance from the listener as the left and right speakers. If the three front speakers were all in the same plane, the sound from the center speaker would reach the listener before the sound from the left and right speakers.

The rear-channel speakers should be placed at 135–150°, as shown in [Fig.14-12](#), and at the same distance from the listener as the front three speakers. This placement isn't always practical, however, so many multichannel products provide a rear-channel delay for those situations in which the listener must sit closer to the rear loudspeakers. Delaying the signals to the rear channels causes the sound from the rear speakers to reach the listener at the same time as sound from the front speakers.

With this array, the front left and rear left speakers can produce phantom images between them along the left sidewall, and the front right and rear right speakers create phantom images along the right sidewall. Correct speaker placement helps to achieve a soundfield that appears to be continuous from front to back, rather than two separate soundfields at the front and rear of the room.

A multichannel music system can also employ seven speakers, with the two additional speakers located at the side positions. This array produces the greatest

spatial realism because the side images are now “hard” rather than phantom, and the feeling of being immersed in a continuous soundfield is enhanced. Multichannel recordings made in a large acoustic and played back through a properly set up 7.1-channel system can be staggering in their ability to create a feeling of actually being inside the original acoustic space. It is even possible to hear discrete sidewall reflections localized along the sidewalls, just as one would hear at the original acoustic event.

Loudspeaker arrays optimized to reproduce film-soundtracks usually employ dipolar, rather than point-source, surround speakers. Dipolar speakers produce sound equally to the front and rear; when they’re positioned on the sidewalls, the listener hears sound from the surround speakers only after it has been reflected from the room’s boundaries. This simulates the array of multiple surround speakers in a movie theater from just two surround speakers, and creates a greater feeling of envelopment. ([Chapter 12](#) includes more detailed information about multichannel loudspeaker placement.)




Fig.14-12 A typical speaker array for multichannel music reproduction employs five identical full-range speakers positioned equidistant from the listener at the locations shown.

Loudspeaker Placement Summary

Loudspeaker placement is the single most important thing you can do to improve your system's sound. It's free, helps develop listening skills, and can make the difference between mediocre and spectacular sound with the same electronics and loudspeakers. Before spending money on upgrading components or acoustic treatments, be sure you've realized your system's potential with correct speaker placement.

After you've found the best speaker positions, install the carpet-piercing spikes (if any) supplied by the speaker manufacturer. Level the spikes so that the loudspeaker doesn't rock: the speaker's weight should be carried by all four (or three) spikes. If you have wood floors that you don't want to mar with spikes, place the round metal discs that are often supplied with the speakers beneath the spikes. To be sure that the spikes are installed at the same depth into the speaker's base, use a bubble level or laser.

You've seen how good speaker placement can give you precise and independent control over various aspects of the sound. You can control both the quantity and the quality of the bass by changing the loudspeakers' distances from the front and sidewalls. The audibility of room resonance modes can be reduced by finding the best spots for the speakers and listening chair. Treble balance can be adjusted by listening height and toe-in. The balance between soundstage focus and spaciousness is easily changed just by toeing-in the speakers. Soundstage depth can be increased by moving the speakers farther out into the room.

I've had the privilege of watching some of the world's greatest speaker designers set up loudspeakers in my listening room for review. At the very highest levels of the art, tiny movements—as little as half an inch—can make the difference between very good and superlative sound. The process can take as little as two hours, or as long as three days. I've often had the experience of thinking the sound was excellent after half a day of moving the speakers around, only to discover that the speakers were capable of much greater performance when perfectly dialed in. Speaker positioning is a powerful tool for achieving the best sound in your listening room, and it doesn't cost a cent. Take advantage of it.

Common Room Problems and How to Treat Them

Treating your listening room can range from simply hanging a rug on a wall to installing specially designed acoustic devices. Large gains in sound quality can be realized just by adding—or moving—common domestic materials such as carpets, area rugs, and drapes. This approach is inexpensive, simple, and often more aesthetically pleasing than installing less familiar acoustic products. You can improve your room with existing household materials, or take the next step of installing dedicated acoustic-control devices. Here are some of the most common room problems, and how to correct them.

1) Untreated parallel surfaces

Perhaps the most common and pernicious of room problem is that of untreated parallel surfaces. If two reflective surfaces face each other, *flutter echo* will occur. Flutter echo is a “pinging” sound that remains after the direct sound has stopped. If you’ve ever been in an empty, uncarpeted house and clapped your hands, you’ve heard flutter echo. Sounding like a ringing that hangs in the air long after the clap has decayed, flutter echo is a periodic repetition caused by the rapid uncontrolled reflection of a sound back and forth between two surfaces. Imagine two mirrors facing each other, the reflections bouncing back and forth between the reflective surfaces creating the illusion of an infinitely receding distance. Flutter echo can blur transient attacks and decays and add a hard, metallic character to the upper midband and treble.

Try clapping your hands in various rooms of the house—particularly the bathroom or hallway. If your listening room has a pinging overhang similar to what you hear in the bathroom, you need to correct this problem.

Flutter echo is easy to prevent. Simply identify the reflective parallel surfaces and put an absorbing or diffusing material on one of them. This will break up the repeated reflections between the surfaces. The material could be a rug hung on a wall, a carpet on the floor (if the flutter echo is between a hard floor and ceiling), drapes over a window, or an acoustically absorbent material applied to a wall. Keep in mind that you need to treat only one of the parallel surfaces to prevent flutter echo. Whatever solution works best for you, killing flutter echo is of utmost importance.

2) Uncontrolled floor and sidewall reflections

It is inevitable that loudspeakers will be placed next to the room's sidewalls and near the floor. Sound from the speakers reaches the listener directly, in addition to being reflected from the sidewalls, floor, and ceiling. Sidewall reflections are the music signal delayed in time, colored in timbre, and spatially positioned at locations different from the source of the direct sound. All of these factors can degrade sound quality. Moreover, floor and sidewall reflections interact with the direct sound to further color the music's tonal character. Fig.14-13 shows how the sound at the listening seat is a combination of direct and reflected sound.

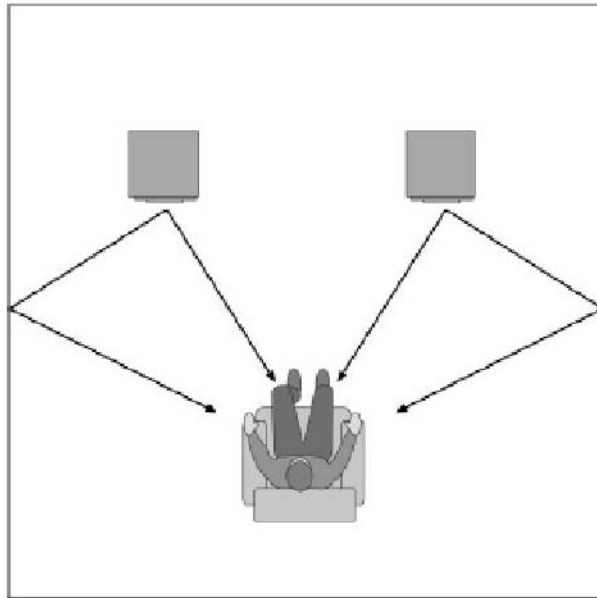


Fig.14-13 The listener receives a combination of direct and reflected sound.

Sidewall reflections color the music's tonal balance in three ways. First, virtually all loudspeakers' *off-axis responses* (the frequency response measured at the side of the speaker) are much less flat (accurate) than their *on-axis responses*. The sound emanating from the speaker sides (the signal that reflects off the sidewall) may have large peaks and dips in its frequency response. When this colored signal is reflected from the sidewall to the listener, we hear this tonal coloration imposed on the music.

Second, the sidewall's acoustic characteristics will further color the reflection. If the wall absorbs high frequencies but not midband energy, the reflection will have even less treble.

Finally, when the direct and reflected sounds combine, the listener hears a combination of the direct sound from the loudspeaker and a slightly delayed version of the sound reflected from the sidewall. The delay is caused by the additional path-length difference between the sound source (the speaker) and the listener. Because sound travels at 1130 feet per second, we can easily calculate the delay time. If the additional difference the reflected sounds travel in Fig.14-13 is 4', the sidewall reflections will arrive at the listening seat 3.5ms (3.5 milliseconds, or 3.5 thousandths of a second) later than the direct sound from the speakers.

The result is a phenomenon called *comb filtering* (see Fig.14-14): a sequence of peaks and notches (note the similarity to a comb) in the frequency response caused by constructive and destructive interference between the direct and reflected sounds. In other words, the phase difference between the two signals causes cancellation at certain frequencies and reinforcement at others, determined by the path-length distance. It all adds up to coloration of the signal at the listening position.

Comb filtering is used as a special effect in the recording studio. Adding a slight delay to a signal and mixing the direct and delayed signals is called *phasing*, characterized by a hollow sound. Darth Vader's voice in the *Star Wars* series undoubtedly includes some phasing. Continuously varying the delay causes the frequencies at which the reinforcements and cancellations occur to sweep the audioband, for an effect called *flanging*. It's called flanging because, in the days before digital-delay lines, the engineer used a tape machine to create the delay, and would put his thumb on the tape reel's flange to vary the delay time and create the characteristic sweeping sound. The most widely heard example of flanging is perhaps the Doobie Brothers' song "Listen to the Music."




Fig.14-14 Comb filtering is a sequence of peaks and notches in the frequency response caused by combining direct and reflected sound. (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988)

The result of these mechanisms—the loudspeakers’ colored off-axis response, the sidewalls’ acoustic properties, and comb filtering—is a sound with a very different tonal balance from that of the direct signal from the speakers. Sidewall reflections are one reason the same speakers sound different in different rooms.

Sidewall reflections not only affect the perceived spectral balance, they also destroy precise image placement within the soundstage. The reflections present “virtual” images of the speakers’ signal that appear on the sidewalls. Although some degree of sidewall reflection adds spaciousness and size, strong reflections increase the *apparent* distance between the speakers. This blurs the spatial distinction between individual images and makes the soundstage less focused and precise. When we hear a center-placed image as partially emanating from positions beyond the outer side panels of the left and right speakers, the tight image focus we seek is destroyed.

Sound is also reflected from the floor and ceiling. Floor reflections tend to cause a reduction in midbass energy, making the sound a little leaner. The ceiling reflection affects the sound less than the sidewall reflections because the soundwaves’ path length is longer. Note that the sound of dipolar speakers, which direct very little energy toward the ceiling, is less affected by ceiling reflections

than is the sound of conventional loudspeakers. Finally, a sloped ceiling is advantageous with conventional speakers when those speakers are placed at the short end of the room. The ceiling slope will tend to direct ceiling reflections away from the listening position.

Reflections reaching the listener from just one room boundary are called *first-order reflections*. These reflections have the greatest amplitude, and have the most effect on the reproduced sound. Fortunately, treating sidewall reflections is simple: just put an absorbing or diffusing material on the side walls between the speakers and the listening position. Drapes are highly effective, particularly those with heavy materials and deep folds. The floor reflection is even easier to deal with: carpet or a heavy area rug on the floor will absorb most of the reflection and reduce its detrimental effects. Low frequencies, however, won't be absorbed by a carpet or rug, leading to a cancellation in the midbass caused by interference between the direct and reflected waves. This is the so-called Allison Effect, named after speaker designer Roy Allison, who first publicized the phenomenon.

Interestingly, the type of carpet between you and the loudspeakers can affect sound quality. Specifically, wool carpet produces a more natural tonal balance than does synthetic carpet. That's because the fibers in wool carpet all have slightly different lengths and thicknesses, which makes them absorb different frequencies. Synthetic carpet, by contrast, is composed of identically sized and shaped fibers that absorb only a very narrow band of frequencies. You can demonstrate this to yourself at a carpet store by speaking into sample pieces of wool and synthetic carpet and listening to the sound of your voice. The wool carpet will produce a more natural timbre.

Sidewall reflections should be diffused (scattered) or absorbed. Diffusion turns the single discrete reflection into many lower-amplitude reflections spread out over time and reflected in different directions (see Fig.14-15).

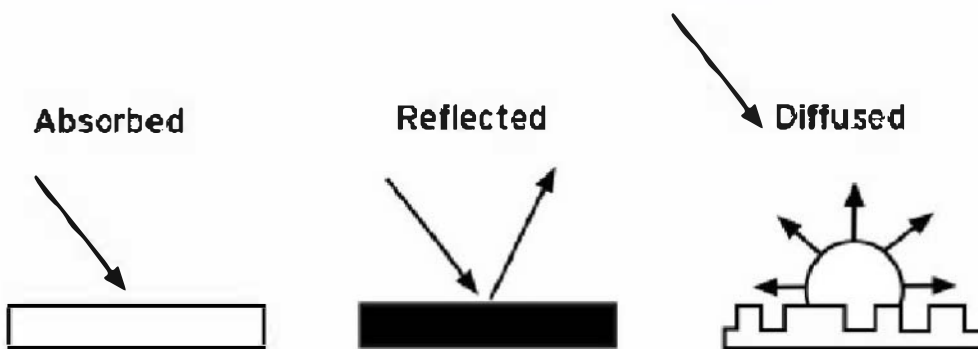


Fig.14-15 Sound striking a surface is absorbed, reflected, diffused, or a combination of the three).

Diffusion can be achieved with specialized acoustic diffusers such as those made by RPG Diffusers (shown in Fig.14-16, and described later in this chapter), or an irregular surface. A open-backed bookcase full of books makes an excellent diffuser, particularly if the books are of different depths, or are arranged with their spines sticking out by different amounts.

The second option is to absorb the sidewall reflection with an acoustically absorbent material. The acoustical foams described later in this chapter will work, but completely absorbing the sidewall reflection with an aggressive foam can make the room sound lifeless and constrict the soundstage's size and spaciousness.

There is some debate in the high-end community about whether sidewall reflections should be absorbed or diffused. Proponents of diffusion argue that the reflected energy is beneficial if converted to many lower-amplitude reflections spread out over time and space, the diffused reflections increasing the sound's spaciousness and air. Absorption proponents suggest that any reflections within the first 20ms of the direct sound degrade the signal from the loudspeakers. Most recording-studio control rooms are designed to provide a "reflection free zone" (RFZ) in which the engineer sits, so that he or she hears only the direct sound from the studio monitors. My experience suggests that absorbing sidewall reflections is better than diffusing them. Diffusing sidewall reflections tends to increase the spaciousness of the sound; absorbing sidewall reflections reduces soundstage width, but results in tighter image focus and a more spatially coherent sound. There is *no* debate, however, about whether or not uncontrolled sidewall reflections degrade a room's sound. They do.



Fig.14-16 An RPG Diffusor (Courtesy RPG Diffusor Systems, Inc.)

An excellent product for controlling sidewall reflections is the Tower Trap, made by the Acoustic Sciences Corporation. This is a tall, cylindrical device with absorptive and reflective (diffusive) sides. Absorption or diffusion can be selected simply by turning the device around. When placed near a sidewall with the reflective side to the back of the room, the absorptive side prevents the first reflections from reaching the listener directly. Some of the energy striking the sidewall is reflected into the Tower Trap's rear (diffusive) side. Most of the sidewall reflection is absorbed, while some is delayed in time, attenuated, and diffused—exactly what we want.

Note that it isn't necessary to treat a listening room's entire sidewall area; the reflections come only from small points along those walls. At mid- and high frequencies, soundwaves behave more like rays of light. We can thus trace sidewall reflections to the listening seat and put the treatments in exactly the right locations. As with light rays, a soundwave's angle of incidence equals its angle of

reflection. That is, the angle at which a soundwave strikes a reflective surface is the same as the angle at which it bounces off that surface.

The technique for tracing side-wall reflections is shown in the series of photographs and illustrations in Figs. 14-17, 14-18, and 14-19. First, mount a strip of reflective Mylar on the side wall between the listener and the loudspeakers. The strip's center should be at the height of your ears when you're sitting in the listening chair. Next, put light sources (two lamps with their shades removed is ideal) where the speakers are normally placed, as shown in Fig.14-17. When sitting in the listening chair, you'll see the two lamps reflected in the Mylar strip (Fig.14-18). The points along the Mylar strip where you see the lamps' bulbs are precisely the points where sound is reflected from the sidewalls to the listening seat. This is where to put the acoustic treatments. The photograph in Fig.14-19 shows how strategically placed acoustical materials (in this case, ASC Tower Traps) can kill sidewall reflections from both speakers. Compare Fig.14-18 to Fig.14-19.

Repeat the process for the left sidewall. If your listening room is symmetrical and the listening position is in the middle of the room, you need use this technique on only one sidewall, then duplicate the acoustic treatment on the other. To maintain acoustical symmetry in the room, both sidewall treatments should be the same.




Fig.14-17 Replace the loudspeakers with light sources and install a Mylar strip.



Fig.14-18 The points of visual reflection are also the points of acoustic reflection.




Fig.14-19 Acoustic treatments placed precisely at the reflection points kill sidewall reflections.

Treating the side-wall reflections from *both* loudspeakers on *each* sidewall improves imaging. The right sidewall will reflect sound from both the right *and* left loudspeakers. The reflection of the left loudspeaker's signal from the right sidewall confuses image placement and constricts soundstage width. This

reflection can be thought of as a kind of “acoustic crosstalk”; we don’t want left-channel information reflecting from the right sidewall into the right ear.

Note that a treatment placed away from the wall affects a larger area of the wall than a treatment directly attached to the wall. The distance between the treatment and the wall causes the treatment to cast an acoustic shadow on the wall, widening the effective absorption area.

This technique can be applied to all reflections in the listening room. If you put a Mylar strip around the entire listening room, any point where you see a reflection of the light bulb where you are sitting or standing is also a sound-reflection point. Additional absorptive or diffusive surfaces can then be placed and oriented exactly where they do the most good. You can also simply sit in the listening position and have a friend move a mirror along a sidewall. The point where you see the speaker’s drivers in the mirror is the point where you should position absorbing or diffusing materials. This technique also works for floor reflections: the point on the floor at which the mirror reflects the tweeter to the sweet spot is the place to add absorbing material.

Even if you don’t go to the trouble of putting up a Mylar strip and lamps, you should do something to treat sidewall reflections. Bookcases, rugs, and drapes are all better than bare walls. But if you really want to get the best from your system, there’s no substitute for professionally designed acoustic treatments.

3) Thick, boomy bass

Thick, boomy bass is a common affliction that can be difficult to control. It often results from room resonance modes, poor speaker placement, poor speakers, or not enough absorption of low frequencies in the listening room. As we will see in the later section on standing waves, the position of your listening-seat can also exacerbate bass bloat.

If thick and boomy bass persists even after you’ve minimized it with careful speaker placement (the most effective method of alleviating the problem), you may want to consider different speakers. If, however, the boominess is minor and you want to keep your speakers, you can make the sound leaner and tighter by adding low-frequency absorbers. These devices soak up low frequencies rather than reflect them back into the room.

Passive low-frequency absorbers simply convert acoustic energy into another form—usually heat within a fibrous material. Low-frequency absorbers can be bought ready-made (such as ASC’s Tube Traps and Tower Traps), built from common materials, or incorporated into an existing room structure. The easiest

and most effective low-frequency absorber for the audiophile is ASC's 16" Full Round Tube Trap. These devices are extremely powerful, soaking up bass below 125Hz—the region below which most domestic materials stop absorbing. A pair of 16" Full Round Tube Traps in the corners behind the speakers significantly tightens up the bass and removes boom and bloat.

Tube Traps employ a plastic sheet beneath the fabric exterior that covers about half the Trap. By rotating the trap, you can adjust the amount of high- and mid-frequency absorption; with the reflective side facing the listener, the Trap acts as a mid- and high-frequency diffusor. With the reflective side facing the wall, the Trap absorbs a broader range of frequencies. This variability is useful when tuning a room; rotating the Trap so that the reflective side faces the listener keeps the sound from becoming overdamped in the midrange and treble. (Tube Trap construction is shown in Fig.14-20.)

US Patent #4,548,292



ACOUSTIC SCIENCES CORP.

Fig.14-20 Tube Trap construction (Courtesy Acoustic Sciences Corporation)

A very inexpensive and effective low-frequency absorber can be made in a few hours for less than \$20. These devices, called *panel absorbers*, have very high

absorption at low frequencies, and can be tuned to the exact frequency and bandwidth required.

Panel absorbers can stand free of or be built right into an existing wall. In a typical built-in absorber, a 4' by 8' frame of 2x4s is nailed on its edge to the existing wall. The inside is sealed with caulking where the wood meets the wall to make it airtight, then filled with fiberglass insulation. Next, a sheet of Masonite or plywood is nailed over the frame. Many tiny holes are then drilled in the panel. Voilà—a low-frequency panel absorber. (This type of perforated panel absorber is also called a Helmholtz resonator.)

Some panel absorbers have no holes, instead using a very thin sheet of material that flexes when struck by sound. The structure's absorption frequency is a function of the airspace depth (2x4s, 2x8s, 2x10s, or 2x12s may be used) and panel thickness. The fiberglass inside the structure widens the frequency range absorbed. By changing the airspace depth and (in perforated absorbers) hole size, panel absorbers can be tuned to any low to mid frequency. Most rooms need broadband absorption in the bass, but panel absorbers can be narrowly tuned to absorb room-resonance problems. A few panel absorbers, tuned to major resonance modes in the low-frequency spectrum, can greatly reduce bass problems in small rooms. To make these rather ugly contraptions palatable in a home, cover them with fabric. (Note that Helmholtz-type resonators shouldn't be covered directly with fabric: they need at least ¼" of airspace to allow the holes to "breathe" freely.)

Freestanding panel absorbers are built the same way, but with a sturdy backing (¾" particleboard, for example) instead of being attached to a wall. Specific details on building panel absorbers—material thickness, perforation size, spacing, etc.—can be found in F. Alton Everest's *The Master Handbook of Acoustics*.

Yet another way to get much-needed low-frequency absorption is build a bass trap into an existing structure—a closet, say. Simply hanging an absorptive material such as acoustical foam or ordinary fiberglass over the closet opening provides low-frequency absorption. This type of structure is called a *quarter-wavelength trap*. The trap will have peak absorption at the frequency whose ¼ wavelength equals the distance between the closet back wall and the absorptive material, and also at odd multiples of the ¼ wavelength. In fact, any absorptive material hung in front of a reflective surface will form a quarter-wavelength trap. Drapes in front of a window will work, but the distance between the window and drapes is so short that the lowest absorption frequency will be in the midband.

Let's calculate the absorption frequency for a closet 2' deep with absorptive material hung at the front. The formula is $F = 1130/4D$, where F is the lowest absorption peak, 1130 is the speed of sound in feet per second, and $4D$ is four times the distance between the reflective closet wall and the absorptive material. The lowest peak absorption frequency for a 2'-deep closet is 141Hz. This structure will also have peak absorption at odd multiples of a quarter wavelength: for instance, at three (423Hz), five (706Hz), and seven (989Hz) times 141Hz, and so on. Drapes hung 6" from a window will provide peak absorption at 565Hz, 1695Hz, 2825Hz, 3955Hz, and so on. Slanting the absorptive material will skew these values along a continuum, making the absorption more even with frequency.

The quarter-wavelength trap is typically limited to providing midbass absorption because of the large dimensions needed to achieve deep-bass absorption. I must stress, however, that the easiest and most effective technique for reducing bass thickness is proper speaker placement. These acoustic treatments should be tried only *after* exhausting every possible option in speaker placement.

4) Reflective objects near the loudspeakers

Reflective objects near the speakers—equipment racks, windows behind the speakers, subwoofers or furniture between the speakers, even power amplifiers on the floor—can cause poor image focus and lack of depth. The best solution is to remove the offending objects. I gained a huge increase in image focus and soundstage depth in my system just by moving my equipment racks from between the speakers to the side of the room. If this is impractical, move the reflective objects as far behind the speakers as possible. Power amplifiers should not extend in front of the speakers' front panels, for example. And for the best musical performance, don't put a big television between the speakers. This is one reason why a music system should ideally be separate from a video system: imaging is degraded by a large piece of acoustically reflective glass near the speakers. Having said that, however, there are a few tricks you can use to minimize the deleterious effect a television has on the reproduction of the music. (These techniques are described in [Chapter 11](#).) Systems that employ a front-projector system and separate retractable screen don't suffer this problem.

If you can't move the offending object, cover it with an absorbing material such as Sonex. Windows behind the loudspeakers should be draped, with the drapes closed when you're listening. There's a large window directly behind my

speakers; I hear a significant increase in soundstage depth with the drapes pulled closed.

A good test of whether imaging is degraded by reflective objects near the loudspeakers is the Listening Environment Diagnostic Recording (LEDR) found on the *Chesky Records Jazz Sampler & Audiophile Test Compact Disc, Vol.1* (Chesky JD37). The test consists of a synthetic percussion instrument processed to sound as if it is moving in an arc between and above the speakers. A system and room with good imaging will make the sound appear to move smoothly and consistently in the intended direction. Any holes in the image will be characterized by the sound jumping from one spot to the next rather than by a smooth, unbroken motion. Such holes could be due to poor loudspeaker quality or placement, but can often be corrected just by moving reflective objects away from the speakers. Best of all, a potentially large improvement in imaging and soundstaging can be had for only a few minutes of your time.

Acoustical Dos and Don'ts

I've summarized this chapter into a few simple guidelines for improving your listening room. This next section is a short course for those readers who don't want to read about absorption coefficients and off-axis response. If you just want some practical tips for getting the most out of your system, this is the section to read. More detail on each of these points can be found throughout the chapter.

1) Loudspeaker placement

Just as real estate agents chant "location, location, location" as the three most important things about the desirability of a house, the three most important ways to improve the sound in your room are loudspeaker placement, loudspeaker placement, and loudspeaker placement. Follow the suggestions in this chapter and spend a few hours moving your speakers around and listening. You'll not only end up with better sound, but become more attuned to sonic differences. All acoustic treatments should be built on a foundation of good speaker placement.

2) Start with good ratios between the room's length, width, and height.

This isn't always possible, but when you're out house-hunting, take your tape measure and the table on good room ratios found later in this chapter. If you can build a room from scratch, or convert a garage or basement into a listening room, choosing optimal dimensional ratios gives you a significant head start in getting great sound from your system.

3) Avoid untreated parallel surfaces.

If you've got bare walls facing each other, you'll have flutter echo. Kill the flutter echo by covering one wall with an absorbent material. Use one of the acoustical foams described, RPG Flutterfree, a rug, or the carpet-like material described later in this chapter. Note that products such as ASC Tube Traps can provide low-frequency absorption, sidewall diffusion, and reduce flutter echo, all at the same time.

4) Absorb or diffuse sidewall and floor reflections.

Bare floors should be covered with carpet between the listening seat and the loudspeakers. Treat the sidewalls between the speakers and the listening position with an absorbing or diffusing material. Avoid having reflective surfaces, such as bare walls and windows, next to the loudspeakers. ASC Tower Traps, the various acoustical foams, and RPG Diffusors are all effective in treating sidewall reflections. Bookcases on the sidewalls also make good diffusers.

5) Keep reflective objects away from loudspeakers.

Equipment racks, power amplifiers, furniture, and other acoustically reflective objects near the speakers will degrade imaging and soundstaging. Move them behind the speakers, or cover them with an absorbing material.

6) Choose a room with a high, sloped ceiling.

This isn't always possible, but a high, sloped ceiling skews the floor-to-ceiling resonance modes and reduces the early reflections at the listening position, resulting in smoother bass and a more spacious and open sound. The speakers should be positioned in the low-ceilinged portion of the room.

7) Balance high-frequency absorbing materials with low-frequency absorbing materials.

Most rooms have lots of high-frequency absorption but little low-frequency absorption. Carpets, drapes, and soft furniture should be complemented with low-frequency absorbers (ASC Tube Traps or home-made panel absorbers) to keep the reverberation time flat across the audioband. Thick, slow bass results from such imbalances in reverberation.

8) Move the listening seat for best low-frequency balance.

Standing waves create stationary areas of high and low pressure in the room. Move the listening chair for best balance. Avoid sitting against the room's rear wall; the sound there will be bass-heavy.

9) Break up standing-wave patterns with irregular surfaces or objects.

Strategically placed furniture or structures help break up standing waves. Large pieces of furniture behind the listening position diffuse waves reflected from the wall behind the listener.

A Short Course in Acoustical Theory

You don't need to know any acoustical theory to get good sound from your system. However, I've included the following survey of basic acoustical principles for those who want to take the next step in understanding how sound behaves in a room.

Listening-Room Resonance Modes

Resonance is the vibration of an object at its natural frequency, determined by the object's material and dimensions. Resonances surround us all the time—from a bell ringing at a certain pitch, to the sound made by blowing across the opening of a soda bottle. A singer breaking a glass with her voice is another example of resonance. It's not that the singer is that loud; when she hits the resonant frequency of the air in the glass, the reinforcement caused by the resonance adds energy to her voice.

As music lovers, we're interested in acoustic resonance: the reinforcement of certain frequencies within an enclosed volume of air, such as a listening room. When excited by the sound from the loudspeakers, the air in the listening room will resonate at particular frequencies, determined by the distances between the room's walls. These resonant points, called *room-resonance modes*, can severely color the bass by creating large peaks and dips in the frequency response. Room-resonance modes impose a sonic signature on the reproduced sound.

The room acts as an equalizer between your loudspeakers and ears, boosting certain frequencies and attenuating others. The result is degraded sound quality. In a poor room, the bass might lack specific pitch definition or sound slow and tubby, certain notes might be more prominent than others, dynamic impact could be reduced—or all of these at once. Resonance modes also cause the room to selectively store energy at the resonant frequencies, and then release that energy more slowly than other frequencies. The result is a tonal blurring and muddying of the sound, along with a distortion of the music's dynamic structure, and a reduction in transient fidelity.

Let's consider what happens along the length of a listening room when the air inside it is excited by sound from the loudspeakers. For now, we'll ignore the room's height and width, and consider only the resonance between the two walls that define the room's length.

A listening room's resonant frequencies are determined by the distance between the room's walls. The farther apart the walls are, the lower the resonant frequency. Specifically, the lowest resonant frequency, called the *fundamental* resonance, occurs when the room's length equals half the wavelength of the sound. Put another way, a resonant mode will occur when the sound's wavelength is twice the length of the room.

Other resonant modes occur at twice this frequency, three times this frequency, and so on. Whenever the length of the room is a multiple of half of the sound's wavelength, a resonant mode will occur. This phenomenon, along with the formula for calculating resonant mode frequencies, is shown in [Fig.14-21](#).

Calculating your listening room's resonant modes is easy. (But remember: we're still considering only the room's length.) The formula is $F1 = 1130/2L$. $F1$ is the first resonant mode, 1130 is the speed of sound in air (in feet per second), and $2L$ is two times the room's length (in feet). If the room is 21' long, its first resonant mode will be 27Hz ($1130/2 \times 21$).

We know that the next mode will occur when the wavelength equals the room's length, at 54Hz ($2 \times F1$), then again at the next multiple of half a wavelength ($1\frac{1}{2}$ wavelengths) at 81Hz ($3 \times F1$), again at 108Hz ($4 \times F1$), and so forth. It is necessary to consider room modes only up to about 300Hz, after which they tend to become so dense that they don't cause problems.

The room's height and width will also create their own resonant modes. If we have an 8' ceiling, the resonant modes will occur at 71Hz ($1130/2 \times 8$), 141Hz, 212Hz, and so on. If the width is 13', the resonant modes will be at 43Hz ($1130/2 \times 13$), 87Hz, 130Hz, 174Hz, 217Hz, 261Hz, etc. What this means to audiophiles wanting the most accurate reproduction possible is that those frequencies will be emphasized in a 21' by 8' by 13' room because of the room's resonance modes. These emphases, or peaks, in the frequency response color the sound. In fact, room-resonance modes can cause bass peaks and dips as large as 30dB. Resonance modes are a significant source of coloration in listening rooms, causing poor bass articulation, boominess, and bass thickness.




Fig.14-21 Resonant modes occur when the distance between the room's walls equals half the wavelength of the sound. Resonant modes also occur at each multiple of half a wavelength.

You can hear (and identify) room-resonance modes most easily by playing the swept sinewave test from the *AVIA Guide to Home Theater* test DVD or other source. This test signal sweeps from a high frequency to a low frequency, accompanied by a readout from the video portion of the disc (when the DVD player is connected to a television) showing the descending frequency of the test signal at any given moment. You can simply listen and hear how some frequencies get much louder than others, or use a sound-pressure level (SPL) meter to measure the effect.

We can see how resonances store energy at certain frequencies, and then release that energy over time, in Fig.14-22. These are actual measurements of my previous listening room, before and after acoustic treatment. The MATT test signal is played through the loudspeakers, a measurement microphone picks up

the signal, and the result is plotted graphically. The MATT signal is a swept sine wave, spanning the range of 20–1950Hz, and interrupted with silence every 1/16 of a second. The signal sounds like a series of blips of ascending frequency, as the signal is on, and then off, for that duration. The measurement result shows how the room modifies the signal, specifically the frequencies at which the room stores energy, and how severely the transient nature of the test signal is distorted. Acoustic energy “hangs” in the room during the silences between blips, filling in the nulls in the chart. The deeper the nulls in the trace, and the flatter the peaks, and the more regular the trace’s shape, the better.

The upper trace looks almost like a squarewave, with deep silences between blips, indicating that the room isn’t storing and releasing much energy. The lower trace, the room’s measured performance before treatment, shows that the nulls are much less deep. This indicates that the room is smearing the energy from the test signal over time, reducing articulation and degrading transient fidelity.




Fig.14-22 Results of the MATT test on an actual room before (lower trace) and after (upper trace) acoustic treatment

Optimizing Dimensional Ratios

We can minimize the frequency-response peaks and dips caused by room-resonance modes by choosing a listening room with dimensional ratios that more evenly spread the resonant modes over the low-frequency band. Let's examine the resonance modes of two rooms of about equal volume. The first room has poor dimensional ratios, while the second room's dimensional ratios are excellent.

The first room is 24' long and 16' wide, with an 8' ceiling. If we calculate and plot the resonant modes for each distance (length, width, height), we end up with a table that looks like this:

Mode	Length	Width	Height
	24'	16'	8'
F1	24Hz	35Hz	71Hz
F2	48Hz	70Hz	142Hz
F3	72Hz	105Hz	213Hz
F4	96Hz	140Hz	284Hz
F5	120Hz	175Hz	355Hz
F6	144Hz	210Hz	
F7	168Hz	245Hz	
F8	192Hz	280Hz	
F9	216Hz	310Hz	
F10	240Hz		
F11	264Hz		
F12	288Hz		
F13	312Hz		

Notice that the third length mode (F3) at 72Hz coincides almost exactly with the second width mode (F2) and the first height mode (F1). These three modes combine, piling up at 70–72Hz to create a huge peak in the response at this frequency. This undesirable situation occurs again at near 140Hz, 213Hz, and 284Hz. The resonant modes coincide because the room's length, width, and height are multiples of each other. Actually, the three modes are perfectly coincident at 70.6Hz, 141.2Hz, etc. I've rounded them so that they are slightly skewed, making the pileup more visible in [Fig.14-23](#) (below). Moreover, there are

no modes between 105Hz and 140Hz, which makes the pileup at 140Hz much more audible. The result will be thick, peaky, and very colored bass.

We can also plot these modes graphically to better visualize room-mode distribution. This graph of mode vs. frequency is shown in [Fig.14-23](#).




Fig.14-23 Plotting the resonant modes graphically

As we can see, the modes are bunched together, with large spaces between them. The gaps between modes are just as detrimental as the coincident modes; they make the adjacent modes much more audible. Moreover, notes falling between modes are reduced in amplitude, making them sound weak and less prominent.

Let's see how this situation can be avoided by a room with better dimensional ratios. Our second room has approximately the same volume, but different distances between its walls. This room is 21' long and 14' wide, with a 10' ceiling. Calculating the resonant modes for each dimension, we see that they are much better distributed:

Mode	Length	Width	Height
	21'	14'	10'
F1	27Hz	40Hz	56Hz
F2	54Hz	81Hz	113Hz
F3	81Hz	121Hz	169Hz
F4	108Hz	161Hz	226Hz
F5	135Hz	201Hz	282Hz
F6	161Hz	242Hz	339Hz
F7	188Hz	282Hz	
F8	215Hz	322Hz	

Fig.14-24 shows these data in graphical form.

We can see by comparing Fig.14-23 with Fig.14-24 that the resonant modes are more evenly distributed in the second room. Although there is some coincidence of modes, they are much less severe than in the first room. This second room will have smoother bass response, less coloration, and a more taut, punchy low end than the first room. Note that larger rooms have inherently more modes, and that the modes will therefore be more smoothly distributed—assuming good dimensional ratios. The smaller the room, the greater the challenge in distributing room-resonance modes. Ideally, the room will have 20 or more axial modes below 325Hz, and the average spacing between those modes will be less than 15Hz. Distribution of modes along the width, length, and height can also be plotted, as in Fig.14-25.




Fig.14-24 Good dimensional ratios better distribute room-resonant modes.

Mode Spacing




Fig.14-25 Another way of graphically representing modal distribution (Courtesy A/V Room-Service, Ltd.)

A sloped ceiling greatly helps in skewing the height modes along a continuum. The resonant modes are naturally spread out rather than occurring at the same frequency, as with a flat ceiling. If a ceiling starts at a height of 8' and slopes up to 12', calculate its modes assuming a height of 10' (the average of the two heights). The sloped ceiling doesn't eliminate resonance modes, but does help to distribute them more evenly. Similarly, splayed walls tend to skew the resonance modes, but don't eliminate them. Splayed walls make calculating room modes more difficult, and don't help much in reducing the audibility of resonance modes.

Room-resonant modes can be made less audible by avoiding square or cubical rooms—i.e., rooms of equal width and height, or width and length, or length and height, or whose measurements are equal in all three dimensions—or rooms whose long dimensions are whole multiples of its short dimensions. If the ceiling is 8' high, we want to avoid another dimension at 8', 16', 24', etc. Note that resonance modes will always occur; good dimensional ratios merely distribute them evenly rather than allowing them to coincide. The table on the next page shows, in descending order of quality, the best dimensional ratios (according to acoustician M.M. Loudon). According to this table, the ideal length and width of a room with an 8' ceiling is 15.2' by 11.2'. These dimensions are calculated by multiplying each dimension by 0.8, because the ceiling is 8' high. For a 12' ceiling, the optimal length and width are 22.8' by 16.8' (multiplying 1.2 by 19' and by 14'). Note that, although these dimensions are different for the three ceiling heights, the *ratio* between the length, width, and height are identical. It's the ratio that determines how smoothly the resonance modes are distributed.

Note also that a larger room with good dimensional ratios will sound better than a smaller room built to those same ratios. The larger room will produce more resonance modes, which can be more closely spaced in frequency and thus

achieve a smoother distribution. Remember that good dimensional ratios don't prevent resonances; they instead distribute those resonances more smoothly over the audioband. Good ratios prevent the coincidence of resonance modes, and the attendant gaps in resonance modes that make the modes more audible.

If you don't have the luxury of choosing the dimensions of your listening room, you can use the techniques described earlier in this chapter to minimize the audibility of room-resonant modes.

Computer software for calculating room-resonance modes is available at a moderate price. For example: Given your room's dimensions, RPG's Room Optimizer software will run a full modal analysis and recommend loudspeaker and listener positions.

This explanation of room-resonance modes and the examples given are greatly simplified. I've just covered what are called *axial modes*—those that exist between one pair of surfaces. Other modes, resulting from two surface pairs (*tangential modes*) and three surface pairs (*oblique modes*, from all six walls of a room), aren't discussed. *The Master Handbook of Acoustics* provides greater depth on this subject.

The 20 Best Dimensional Ratios (according to M.M. Loudon)

Quality	Height	X	Y
1	1	1.9	1.4
2	1	1.9	1.3
3	1	1.5	2.1
4	1	1.5	2.2
5	1	1.2	1.5
6	1	1.4	2.1
7	1	1.1	1.4
8	1	1.8	1.4
9	1	1.6	2.1
10	1	1.2	1.4
11	1	1.6	1.2
12	1	1.6	2.3
13	1	1.6	2.2
14	1	1.8	1.3
15	1	1.1	1.5

16	1	1.6	2.4
17	1	1.6	1.3
18	1	1.9	1.5
19	1	1.1	1.6
20	1	1.3	1.7

(X can be length and Y width, or vice versa)

Standing Waves

If you've seen a cup of coffee sitting on a vibrating surface, you've seen standing waves. The vibration excites the liquid, causing waves to spread out from the center and reflect or bounce off of the cup's inner rim. At some points, the waves reinforce each other; at other points, the waves cancel. The result is a stationary pattern of peaks and dips in the surface of the coffee. Although the two waves are in motion, they create a stationary wave formed by the interaction of the two moving waves. This is exactly what happens to sound in a listening room.

Standing waves are stationary areas of high and low sound pressure in a room. They are caused by constructive and destructive interferences between the incident (direct) sound and the reflected sound from a room boundary. For example, when the reflected waveform's positive pressure phase is superimposed on the incident waveform's positive pressure phase, the two waves reinforce each other (constructive interference), producing a *peak*. If the reflected wave's negative pressure phase meets the incident wave's positive pressure phase, the two waves cancel each other (destructive interference), producing a dip or *null*. These interactions produce a specific pattern of stationary areas of more and less bass for each resonant frequency in the room.

Let's look at how a standing wave is set up along the length of a 20'-long listening room. If we excite the air in this room at its lowest resonant frequency, the pressure maxima will occur at the room walls, the pressure minima in the center. (See Fig.14-26. This drawing is analogous to Fig.14-21, which showed how a resonant mode occurs when half a wavelength equals the room's length.)

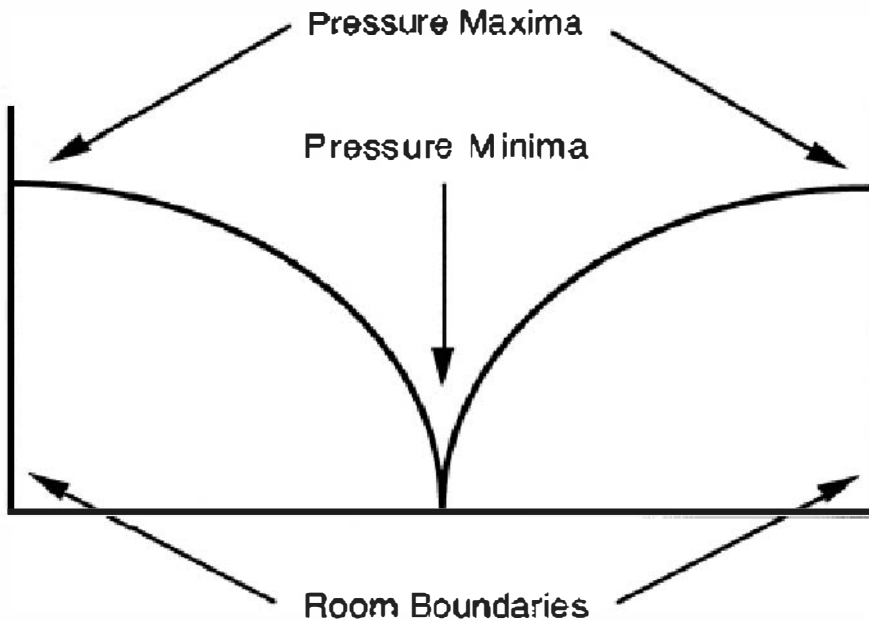


Fig.14-26 Maximum pressure occurs at room boundaries

By playing a tone at the room's resonant frequency and walking around in the room, we can hear the sound get louder and softer. In actual listening rooms, things aren't as neat as in Fig.4-26: the effects of other room boundaries and objects (furniture) in the room will skew this theoretical representation. In addition, the associated resonant modes at multiples of the modal frequency tend to make the pattern of standing waves far more complex.

Nevertheless, standing waves create areas in the room in which the bass sounds heavier, and some in which it sounds thinner. Play music with continuous bass energy (pipe organ music works well) and move your listening chair forward and backward; the sound will be heavier in some areas, leaner in others.

Although this unequal pressure distribution is undesirable, we can use standing waves to our advantage. If your system tends to sound thick and heavy in the bass, try moving your listening chair forward or backward until the bass is a little smoother and leaner. Or use the same technique to add fullness and weight to thin-sounding speakers.

Reverberation

The sound in a room can be thought of as a combination of three components: 1) direct sound from the loudspeakers; 2) discrete early reflections from the floor, ceiling, and sidewalls; and 3) later, more diffuse reflections, called reverberation. This concept of sound having three components is illustrated in Fig.14-27:




Fig.14-27 Sound in a room is composed of direct sound, early reflections, and reverberation.

Reverberation is such a fundamental part of sound that we take it for granted. Although we don't hear it as a separate component of sound, reverberation adds warmth and space to music. If reverberation were completely missing from a sound, we would immediately identify the sound as being unnatural. You can hear how strange the absence of reverberation is on the *Denon Anechoic Orchestral Recording* (Denon PG 6006). The orchestra, recorded in an anechoic chamber (a room in which all sound reflections are completely absorbed), sounds completely different from even a recording made in a small, "dry" concert hall.

Reverberation in a room takes a finite time to decay into inaudibility. The time it takes for sound to decay by 60dB is called the room's RT_{60} (RT stands for "reverberation time"). A small, acoustically dead rock recording studio may have

an RT_{60} of 0.1 second; a cathedral may have an RT_{60} of six or seven seconds. A room's RT_{60} quantifies how acoustically “live” or “dead” the room is. In practice, sound in a room doesn't decay the full 60dB, because the decaying sound becomes swamped by the room's ambient noise.

Although the concept of reverberation is less useful in the treatment of small spaces such as hi-fi listening rooms than it is for concert halls, it is nonetheless helpful in balancing sound-absorbing materials in the room. A room with reflective surfaces will have a longer reverberation time than a room with absorptive surfaces such as carpets, soft furniture, and drapes. These materials absorb sound rather than reflect it back into the room. You want to choose materials in the listening room that contribute to an optimal RT_{60} , and balance those materials so that the reverberation time is constant over the audioband.

But what is the ideal reverberation time in a music listening room? And what combination of reflective and absorptive materials will produce the optimal reverberation time?

The quick answer is that optimal reverberation time in a listening room varies with the room's volume. For a 3000-cubic-foot room, a good reverberation time is about 0.4 second. For a 20,000-cubic-foot room, about 1 second is ideal. A surprisingly accurate measure of a room's reverberation time can be made with a handclap or popped balloon and a stopwatch. Have someone clap his hands loudly, or pop a balloon. Start the stopwatch when you hear the sound, and stop it when the sound has completely died away. Repeat the measurement several times and average the results. The average, in seconds, is a fairly reliable indicator of a room's RT_{60} .

Unfortunately, an optimal overall reverberation time isn't enough: we must ensure that the room's reverberation time is roughly the same at all audio frequencies. A room with thick carpet and heavy draperies—materials that absorb high frequencies, yet do nothing to low frequencies—will have a longer reverberation time at low than at high frequencies. This condition can make a hi-fi system sound slow and thick in the bass, or dead in the treble. Reverberation time is thus specified at six frequencies: 125Hz, 250Hz, 500Hz, 1kHz, 2kHz, and 4kHz. Our goal is to make the reverberation time roughly equal at all of these frequencies by choosing the appropriate mix of acoustical treatments. We do this by mixing different materials, each having absorption characteristics that vary with frequency. Let's look at how these surfaces affect the sounds striking them.

Every surface in a listening room will absorb, reflect, and diffuse (scatter) sound in different proportions depending on the surface material and the audio

frequency striking the surface. We'll deal first with absorption and reflection. The degree to which a surface absorbs sound is called its *absorption coefficient*. This is a measure of the percentage of sound it absorbs, specified at the six frequencies listed in the previous paragraph. An absorption coefficient of 1.0 means that 100% of the sound is absorbed. This condition is likened to an open window—no energy is reflected back into the room. An absorption coefficient of 0.1 means that 10% of the sound is absorbed and 90% is reflected back into the room. Absorption coefficients of more than 1.0 are possible with a material that presents a surface area to the soundfield that is greater than its area of wall contact, such as some wedge-shaped foams.

The absorption coefficients for two common materials—heavy carpet over felt, and drywall on studs (½” thick on 2x4s, spaced 16” on center)—are as follows:

	125Hz	250Hz	500Hz	1kHz	2kHz	4kHz
Carpet	0.08	0.24	0.57	0.69	0.71	0.73
Drywall	0.29	0.10	0.05	0.04	0.07	0.09

We can quickly see that carpet absorbs virtually no energy in the bass, but a significant percentage of the energy in the treble—nearly 75% at 4kHz. Conversely, drywall absorbs (and leaks out of the room) a moderate amount of energy in the bass (29% at 125Hz), and reflects virtually all midband and treble energy. Drywall on studs absorbs low frequencies by diaphragmatic action: sound striking the drywall causes it to flex. This flexing converts mechanical motion into a minute amount of heat.

Now, suppose we built a room completely covered in thick carpeting—like a rock band's garage studio. The carpet would absorb nearly all the treble and reflect nearly all the bass. The result would be a very short reverberation time at high frequencies and a very long reverberation time at low frequencies. The room would sound thick, heavy, and congested in the bass, and generally unpleasant. Moreover, the room would store energy in the bass through resonance and the release of that energy over time, smearing transient signals. A short reverberation time in the midband and treble, accompanied by boomy bass, causes a loss of intelligibility of sung words. In such a room, each band member keeps turning up his volume so that he can hear himself over the other members. But turning up the volume only makes the problem worse.

At the other extreme, a room of bare drywall, a tile floor, and no high-frequency-absorbing materials would sound bright, hard, and thin. The trick is to


balance absorptive materials so that the amount of absorption is roughly the same at all six frequencies. In practice, it's difficult to get enough low-frequency absorption; most rooms have longer reverberation times in the bass. Fortunately, a slightly longer reverberation time at low frequencies is desirable to give the sound warmth.

Reducing a room's reverberation time by adding absorbing materials has the added benefit of broadening the resonances, making them less audible.

Fig.14-28 shows a computer model of my former listening room's reverberation time. The upper plot is the reverberation time before installation of a complete room-acoustics package. The two lower plots are the ideal reverberation time, and the actual reverberation time after room treatment.

Room Length = 21 ft. 0.6 in.
 Room Width = 14 ft. 6.0 in.
 Ceiling Height = 9 ft. in.

Room Volume = 2746 cf.
 Wall Area = 640 sf.
 Ceiling/Floor Area = 305 sf.



	31	63	125	250	500	1000	2000	4000	8000	Average	
W/O	2.81	1.61	1.13	0.97	1.30	0.90	0.72	0.58	0.72	0.93	sec
With	0.85	0.49	0.34	0.32	0.36	0.32	0.29	0.28	0.36	0.32	sec
Ideal	0.7	0.5	0.35	0.35	0.35	0.35	0.35	0.35	0.35	0.35	sec
Room Mode Bandwidths W/O ARS Treatment .										2.4	Hz.
Room Mode Bandwidths with ARS Treatment										6.9	Hz.

Fig.14-28 Computer model of my former listening room showing reverberation time without treatment, with treatment, and the ideal (Courtesy A/V RoomService, Ltd.)

You can calculate your room's reverberation time. First, multiply the absorption coefficient of every surface in the room by the square footage of that surface at the six frequencies. For example, if you have 294 square feet of carpet, multiply 294×0.73 (absorption coefficient of carpet at 4kHz) to get 215

absorption units. These absorption units are called **sabines**, after Wallace Clement Sabine, the American physicist who founded the science of architectural acoustics. Boston Symphony Hall, designed by Sabine in 1900, was the first concert hall to be designed using the new field of architectural acoustics.

Multiplying each surface area's square footage by the absorption coefficients at the six frequencies gives us a table of every material in the room and how many sabines of absorption that material provides at the six frequencies: 125Hz, 250Hz, 500Hz, 1kHz, 2kHz, and 4kHz.

To convert the total sabines at each frequency into a reverberation time at that frequency, use this formula:

$$RT_{60} = \frac{(0.049)(V)}{(S) (a)}$$

where

RT_{60} = reverberation time, seconds

V = volume of room, cu ft

S = total surface area of room, sq ft

a = average sabine absorption coefficient

This method of calculating reverberation time works best with large rooms, and assumes that the materials are uniformly distributed throughout the room. Nevertheless it will give you an idea of the balance between low- and high-frequency absorption. If you perform these calculations on your own room and find a large difference in the total sabines between 125Hz and 4kHz, you can add low-frequency absorbers to even out the reverberation time. The acoustical materials described earlier in this chapter can be used to achieve a smooth reverberation time with frequency. A full list of materials and their absorption coefficients is included in *The Master Handbook of Acoustics*, mentioned earlier in this chapter.

Building a Listening Room from Scratch

This chapter cannot guide someone through the long and complicated process of building a dedicated listening room from the ground up. Instead, I recommend two approaches: The first is to study the subject and design the room yourself, bringing in an acoustician for consultation two or three times during the design process. You can end up with a good room without spending a fortune. The acoustician can be a freelance contractor, or be associated with a company that sells acoustic treatments.

The second approach is to commission a full-fledged design from a design firm that specializes in architectural acoustics. This way is ideal for those with large budgets or no time to learn acoustic design themselves. The design firm will often oversee construction, providing you with a completely finished room. This can be a *very* expensive proposition. A custom room built with this approach is shown in [Fig.14-29](#) (photo courtesy Pure Audio, Scotts Valley, California; www.balabo.com).



Fig.14-29

Isolating the Listening Room

Many listeners want to play music at high levels without disturbing neighbors or other family members, or to listen to music in a quiet environment without outside sounds intruding on the experience. This brings us to a specialized area of acoustics called isolation. Acoustic isolation prevents sound in one area from escaping into another area. A good example of successful isolation is the midtown Manhattan recording studio that records during business hours.

Keeping sound in the listening room from getting into the rest of the house is a daunting challenge. It requires good design, precision construction, and lots of money. There are, however, a few principles and techniques that can help minimize sound transmission.

Very briefly, some of the sound escapes from the listening room into the rest of the house through doors and other openings. Even small cracks, such as the space between the floor and door, let through large amounts of acoustic energy in a phenomenon called flanking. To reduce flanking, use a solid-core door and seal its perimeter with weather-stripping or refrigerator seals. You can also buy acoustic doors, complete with jambs and seals, from acoustic-supply houses. Look for the ads in the back of *Mix* magazine or other professional audio publications.

Second, isolating a room requires lots of mass between the sound source and the area to be protected from sound. Thick walls, double walls with unbridged air cavities, lead sheeting, and double solid-core doors are all fundamental prerequisites to achieving high isolation. Most audiophiles don't need such isolation, but if you do, it's very expensive. Halfhearted attempts at isolation just don't work; it takes a serious effort to achieve adequate isolation.

The effectiveness of a wall's ability to reduce the amount of sound it lets through is expressed as its Sound Transmission Class (STC). For example, a standard wall built with 2x4 studs and Sheetrock has an STC of 34dB without insulation (as you would typically find in a home's interior walls), and 36dB with 3.5" of fiberglass insulation. Using a 2x6 bottom plate with 2x4 staggered studs (studs alternating on each side of the wall so that the studs don't touch both pieces of Sheetrock) increases the STC to as high as 52dB (with insulation). A full double 2x4 wall with an unbridged air cavity has an STC of 58dB. As you can see, the more elaborate (and expensive) the construction method, the greater the wall's ability to attenuate sound passing through it.

A related subject is the room's *noise criterion* (NC), a numerical rating of the level of background noise present in the listening room. The lower the NC rating (NC20 is excellent, for example), the lower the room's ambient noise. A quiet room makes low-level detail more audible, contributes to the sense of space (reverberation decays are better resolved), and fosters the impression of a "black" background from which the music emerges.

The Walldamp Technique

The modern Western construction technique of drywall over a frame of wooden studs introduces audible problems not found in parts of the world that build with thicker, denser materials. Specifically, sound striking a drywall-over-stud wall causes the wall to resonate at about 70Hz. Sound is converted to mechanical energy (the wall moving), then back to acoustic energy (the vibrating wall acts just like a loudspeaker's diaphragm). The wall hums away along with the music, introducing a drone-like sound of a specific frequency. What's worse, the acoustic energy produced by the wall is released relatively slowly over time, altering music's dynamic characteristics.

This phenomenon, called *wall-stud resonance*, introduces a specific coloration to the reproduced sound. You can demonstrate this by hitting a drywall-over-stud wall with the ball of your fist and listening to the resulting thud. The sound produced by the motion of the wall has virtually the same character, whether the initial sound was a double bass or an explosion in a film soundtrack. Consequently, the music is overlaid with a common signature that degrades fidelity. Homes built with thick plaster walls or concrete block don't suffer from this problem.

While acoustic absorbers can stabilize the air in a listening room, they do nothing to address wall-stud resonance. Keeping the room's walls stabilized under dynamic conditions (loud transient signals) requires careful attention to wall construction. If you're building or remodeling a home, you can greatly reduce wall-stud resonance and isolate the listening room at the same time. One method is to buy dual-layer drywall that has damping material between the layers. Unfortunately, this is a specialized material that must be purchased by the truckload.

The second option is ASC's Walldamp material and construction technique, shown in Fig.14-30. The first step is to install a resilient metal channel between the studs and the first layer of drywall. The resilient channel acts as a spring, isolating the wall from the studs and preventing the 70Hz wall-stud resonance from occurring. If left like this, however, the room would sound horrible; the entire wall would act as a drumhead, resonating whenever acoustic energy struck it.

This is where Walldamp and the second layer of drywall come in. The Walldamp in Fig.4-30 consists of squares of tarpaper coated on both sides with a special glue. When installed between two layers of drywall, the glue absorbs

energy by a mechanism called *shear stress damping*. That is, when the two pieces of drywall try to move relative to each other, the glue absorbs the energy on a molecular level. Specifically, the glue converts the mechanical energy of the wall's motion into a tiny amount of heat. Consequently, the acoustic energy striking the wall is absorbed by the damping action of the Walldamp instead of being reflected back into the room. Rather than shuddering when struck by sound, the wall behaves more like a thick plaster wall.




Fig.14-30 ASC's Walldamp technique damps wall vibration and increases isolation. (Courtesy Acoustic Sciences Corporation)

In addition, the wall now acts as a large bass absorber at very low frequencies —just what most listening rooms need. Fewer freestanding bass absorbers are needed because now the entire *room* is acting as a bass absorber. Another benefit of this technique is that less of the sound inside a listening room escapes into the rest of the house: You can play music louder and later, without disturbing family

members. Best of all, this construction technique adds very little to the cost of building or remodeling.

DSP Room Correction

We've seen in this chapter that the listening room is a crucial link in the music playback chain. Every listening room, no matter how good, radically alters the sound produced by a pair of speakers. The room acts as an equalizer, boosting some frequencies and attenuating others. This change in frequency response can be as much as 30dB in the bass, adding large amounts of coloration to the music. Moreover, the room briefly stores and releases energy at certain frequencies, causing some bass notes to "hang" in the room longer. The result is a thickness and lack of clarity in the bass, along with a reduced sense of dynamic drive.

It's been said that a loudspeaker designer has 100% control over his product's sound above 700Hz, 50% control from 300Hz to 700Hz, and only 20% control below 300Hz. This aphorism reflects the increasing influence of the listening room on sound quality at lower and lower frequencies.

Although the listening room imposes far more severe colorations than any electronic component, we've grown to accept this musical degradation as an inevitable part of music listening. Even custom-built, acoustically designed rooms housing perfectly positioned loudspeakers will introduce significant colorations. All rooms introduce colorations: the better ones simply do less damage. Attempts to reduce room-induced colorations with analog equalizers have met with limited success. Analog equalizers add their own audible problems, and are limited in resolution. Digital equalizers can have greater resolution than analog ones, but are still a crude tool for dealing with room problems.

The advent of digital signal processing (DSP) provides a powerful new tool for removing the effects of the listening room on reproduced music. DSP is the manipulation of audio signals by performing mathematical calculations on the numbers that represent the music. DSP chips can be programmed to perform as filters or equalizers with many hundreds of frequency bands, and with characteristics impossible to realize with analog equalizers. The trick is how to put this advanced technology to work to surgically correct room-induced colorations.

Far more sophisticated than a digital equalizer, a DSP room-correction system analyzes the response of your loudspeaker/room combination with high resolution, then creates digital filters to remove the room-induced problems. The room-correction system not only removes the room's effects, it also corrects for speaker colorations. If your speakers have a peak of energy at 2kHz, for example,

the room-correction system can remove this coloration. In addition, the room-correction system will make the speaker/room response identical for each speaker. This technique removes the room's effects on the recorded sound electronically, *before* it is reproduced by the speakers. Put another way, the room-correction system distorts the signal in a way that counteracts the distortion imposed by the listening room and speakers. The result is flat response at the listening position.

Here's how room correction works: You (or your dealer) position the supplied measurement microphone at the listening seat, connect the room-correction system hardware to a personal computer, and run the supplied software on your PC. The room-correction system hardware outputs a series of pulses that are reproduced as sounds by your loudspeakers. The sounds of those pulses in your room are then analyzed by the software running on your PC. The software then configures filters in the DSP chips that form an inverse of your room's response. Peaks in the measured response become dips in the filter, and vice versa. The filter response is a mirror image of your room's response. Newer, more sophisticated room-correction systems eliminate the need for the PC, and have greatly simplified the user interface. Through its inclusion in home-theater receivers, DSP room correction has become a mainstream technology.

The filter parameters calculated by the software are downloaded into the room-correction system's memory during this setup routine. The room-correction system hardware stays in your rack and performs the filtering on all signals running through it. If you change speakers, speaker positions, or listening spot, you'll need to remeasure the room to create new filter coefficients (about a 15-minute job, once you've been through it a few times). The filter parameters are unique to your speakers, room, speaker positions, and listening position.

At first glance, you might think that the room-correction system should produce perfectly flat response at the listening position. In practice, a perfectly flat response sounds thin, bright, and bass-shy. The ideal response has a rising response in the bass and a falling response in the treble. Some room-correction systems don't give you the option of choosing the curve; you're stuck with whatever the designer thought was ideal. Such systems often improve the bass performance, but add excessive brightness to the overall sound. Other DSP room-correction systems allow you to choose one of several "target curves," or even to create your own curve. This target curve will be the response at the listening position.

An important technical factor in a room-correction system is its *frequency resolution*. This specification indicates the correction system's ability to correct

very narrow peaks and dips in the frequency response. For example, a room-correction system with a frequency resolution of 10Hz wouldn't be able to correct a peak or dip in your room that was 8Hz wide. Some room-correction systems have a frequency resolution of less than 1Hz, which is fine enough to follow any peaks and dips.

Note that because room-correction systems operate in the digital domain, analog signals from an LP front end must be digitized, then converted back to analog. That's why most room-correction systems include an integral analog-to-digital converter. Digital signals, such as from a music server, are input directly to the room-correction system, which either converts that digital signal to analog for driving a power amplifier, or outputs a digital signal for conversion to analog with an outboard digital processor. Although room correction can produce profound improvements in bass quality and image focus, some audiophiles are reluctant to digitize their LPs. The purity of an analog source is inevitably degraded by the digital conversion process. You must consider, however, whether the loss of quality from digitizing an analog signal is a greater or lesser evil than the degradation imposed by the listening room and speakers.

My experience with room correction suggests that the technology is extremely effective in removing bass colorations and restoring clarity to reproduced music. With room correction, the sound has a "lightness," agility, and clarity that are difficult to achieve in any other way. The overall sound seems to have less midbass boxiness, and extreme low frequencies have greater dynamic impact. In addition, soundstaging improves because identical response from both loudspeakers is essential to precise image focus. (That's why each speaker must be toed-in by exactly the same amount.) Because the response of each speaker/room is different, without room correction image size and placement can change as an instrument moves between musical registers. With identical output from each speaker, image focus becomes tighter and the soundstage is better delineated.

On the downside, DSP room correction can radically alter your system's sound. If you like your system's overall tonal balance, adding a room-correction system can be disappointing. Although room correction nearly always results in better bass, it often comes at the price of an unnatural midrange and treble balance.

If your loudspeakers are of the subwoofer/satellite variety, you can use room correction to improve the bass performance without affecting the midrange and treble. Room-correction devices are available that operate only on the signal driving the subwoofer, leaving the signal driving the satellite speakers unaltered. This approach gives you the best of both worlds.

15 System Setup Secrets, Part Two: Expert Tuning Techniques, Accessories, Equipment Racks, and AC Power Conditioners

Introduction

I'm privileged to be a continuing student of the art of audio-system setup in the best classroom imaginable: As a full-time audio writer since 1989, I've watched the designers of the world's top high-end equipment install their products in my home for review. As you can imagine, these designers are highly motivated to get the best possible sound from their products. I've watched these industry veterans pull from their bags of tricks all of the techniques and insights for optimizing a system's performance that they've gained over lifetimes in the field.

I realized early on the unique position I was in, and dedicated myself to learning these setup techniques. Although there are many well-known, standard setup procedures, many of these elite designers have developed their own methods for getting maximum performance from an audio system. It's been fascinating to watch these masters in action, and to listen to the results of each change. Frequently, a system that sounds quite good after an hour of tweaking can sound spectacular after another couple of hours of experimentation.

I'll tell you about just one such setup from early in my reviewing career that vividly demonstrated the power of system tweaking. The review was of a pair of loudspeakers that had been shipped to me a few days before the designer was scheduled to visit. I was eager to hear the speakers, so I uncrated and set them up

myself the day before the scheduled visit. I spent the evening listening to the speakers, and thought they sounded excellent.

The next day the designer arrived and spent about four hours adjusting the speaker positions by listening and then moving the speakers, repeating this process again and again. At the end of the day he wasn't satisfied, so he planned to return the next morning. After he left I spent the evening listening, and was astounded by how much better the system sounded. The improvement was so profound that I wondered why the designer felt the need to return the next day. But an additional four hours of adjustment the following day resulted in another wholesale improvement in their sound. The designer, who knew the speakers' performance potential, made exquisitely fine adjustments to their positions and toe-in angle. He tightened all the driver-mounting bolts, cleaned all the RCA plugs and jacks, applied contact enhancer to the blades of the AC cords, and performed myriad other small tweaks that, cumulatively, made a significant difference. He pronounced the system fully dialed in and headed for the airport. I spent a third consecutive evening listening, with the previous two evenings fresh in my mind. The transformation was staggering. It was hard to believe I was listening to the same source components, same amplification, same interconnects and speaker cables, and same loudspeakers. I would never have thought it possible that the system had this hidden potential, waiting to be unlocked through the application of skill and effort.

In this chapter I'll share with you the best of what I've learned. You can use a few of the techniques to improve your own system, or study this chapter and practice the craft to become a master of the art of system setup.

I've divided the chapter into three sections. The first describes the wide range of accessories available to the audiophile. The second details specific procedures for improving a system's sound. Some of these procedures require buying an accessory (usually inexpensive), and some are techniques that require nothing more than your time, patience, and listening skill.

The third section goes into some depth about AC power conditioners, AC cords, equipment racks, and vibration isolation. Although these products are traditionally categorized as accessories, AC conditioners and equipment racks are fully equal in importance to any other component. In fact, good AC power and a vibration-resistant rack should be considered the platform on which to build the rest of your system.

Nonetheless, I can't overstate the need to base the tweaks described in this chapter on a solid foundation of correct speaker placement. Correctly positioning

your speakers is the single most important thing you can do to get the best sound from your system. Only after you've correctly positioned your speakers to avoid major problems created by the interaction of the speakers with the room should you begin the process of fine tuning described here. If the speakers and listener aren't positioned in the right places, none of the other setup techniques will matter. In fact, speaker placement and room acoustics are so important that I devoted the entire previous chapter to them.

Accessories

The vast array and wide availability of audio accessories bear witness to audiophiles' need to squeeze out those last bits of musical performance from their systems. Accessories can not only make your system sound better—they're fun to try. It's rewarding to discover some simple trick or product that improves your system's sound at little or no cost. Some of the improvements rendered by accessories are marginal, but when applied in concert with other tweaks, the cumulative effect can become significant. This section surveys available accessories, describes what they do, and explains how they work. Note that I don't endorse every accessory described here—I simply report on what's available.

How to Choose Accessories

Some of the accessories described in this section can make your system sound better than you thought it could—and at a modest cost. Other accessories can not only fail to improve the sound, but can actually degrade it. To make matters even more complicated, many accessories are completely worthless—nothing more than snake oil sold by less-than-honest promoters that have no effect on the sound. Finally, the effectiveness of any accessory can vary from system to system.

Fortunately, there's a simple, foolproof way of deciding which accessories are worth your money and which aren't: Listen to the accessory in your system before you buy. Most dealers will let you borrow the accessory overnight, or return it for a refund if you don't hear an improvement. If you hear an improvement that seems commensurate with the product's price, buy it. If you don't, return the product. It's that simple. And if the accessory is inexpensive but not returnable, you can try it at minimal cost. Some CD treatments cost only a few dollars; if they don't work, they haven't cost you much. But by all means, try some of the devices described in this section. They can often make the difference between merely good and superlative sound.

Any pursuit will attract some followers who have abandoned reason, and high-end audio has more than its share of nonsensical, even bizarre accessories. Most of these products are worthless. Examples include an alligator clip, to be attached to a speaker cable to prevent the speaker's "gravitational influence" from affecting the audio signal. Another claims to "energize" the electrons in an audio system with a combination of lithium salts and cobra venom.

We must be careful, however, in branding all odd-sounding products as fraudulent. Many accessories whose value is now not in question—AC power cords and vibration-control devices, for examples—were once dismissed as worthless by those who had never listened to their effects. In the early days of CD, it was assumed that all CD transports and digital interconnects must sound the same because they all produce the same binary datastream of ones and zeros. To suggest otherwise was the height of audiophile lunacy. A more contemporary example is reports of sonic differences between uncompressed WAV and losslessly compressed FLAC files of the same recording; when FLAC files are decompressed, they are bit-for-bit identical with the WAV file, yet they sound different. It was only after much critical listening and closer scrutiny that other mechanisms for introducing differences in the sounds of WAV and FLAC files, and between digital interconnects, were discovered and scientifically

documented. Today it is undisputed that different digital interconnects introduce varying amounts of jitter that produce an analog-like variability in their sound quality. Similarly, decompressing a FLAC file on the fly—i.e., without using playback software that buffers the data—degrades the sound. We must listen and trust our ears, and not rely solely on accepted orthodoxy, no matter how elegant or logical that orthodoxy may appear. It's the listening experience that points the way toward the most fruitful paths of scientific inquiry.

The bottom line with any accessory, as with any audio component, should by now be familiar: Listen before you buy. If the device makes an audible improvement, buy it—regardless of whether or not there's a convincing scientific explanation for what you hear.

Cable Enhancers

A cable enhancer is a device that breaks in line-level interconnects and loudspeaker cables faster and more completely than running music through them. Both ends of the interconnect or cable are connected to the enhancer, which puts out a high-level signal to produce current flow in the cable. Cable enhancers cost about \$200. Many dealers, however, will break in your new cables on an enhancer at no charge. Note that the improved sound of a cable after break-in is largely due to charging the cable's dielectric (the insulating material around the conductors), a charge that dissipates over time. It may thus be necessary to periodically "re-enhance" your cables, especially if they haven't been used in some time.

RF Filters

RF stands for “radio frequency” and describes unwanted noise at frequencies used for radio transmission that can contaminate audio signals. An RF filter is a small, cylindrical device that fits around a cable or AC cord to prevent high-frequency noise from passing through the cable to the audio component connected to it. The filter is, in essence, a choke that turns the cable or power cord into a low-pass filter. That is, the cable or power cord with an RF filter around it allows the relatively low frequencies of the audioband to pass, but blocks the very high radio frequencies. RF filters can be put around interconnects at the preamp end to prevent RF from getting into the preamp, around interconnects at the power amplifier, or around AC cords. They cost about \$25/pair, are easy to install and remove, and can provide an inexpensive improvement in your system’s sound. I’ve had mixed results with RF filters; in my view, a better solution is well-designed AC power cords and interconnects.

Burn-in Signals

Available on some setup CDs or as downloadable files, these special signals reportedly improve a system's sound by demagnetizing the circuits inside the components. There's no scientific basis for the efficacy of burn-in signals, but they do seem to work.

Contact Cleaners

Contact cleaners remove the residue and oxidation that build up on RCA plugs and jacks, interconnect terminations, and loudspeaker and power-amplifier terminals. The fluid comes in a small bottle, and is usually supplied with pipe cleaners or a brush for applying the liquid. Some cleaners come in spray form, but the contacts should still be rubbed to remove oxidation. I recommend DeoxIT, an inexpensive and effective product from Caig Laboratories. The procedure for cleaning contacts and vacuum-tube pins and sockets is described later in this chapter.

Contact Enhancers

A contact enhancer is a fluid applied to the electrical contacts in your system. Contact enhancers promote better conductivity between a plug and jack, thus improving the sound. They work especially well on phono-cartridge and vacuum-tube pins, and are also effective on the blades of AC power cords (Fig.15-1). Contact cleaners and contact enhancers unquestionably can improve a system's sound.



Fig.15-1 Applying contact enhancer to the blades of AC cords is a simple and effective tweak.

Tube Dampers

A tube damper is a ring-like device that fits over a vacuum tube to reduce its vibration. Tubes are microphonic—they can convert physical vibrations into an electrical signal that is then mixed in with the musical signal. These vibrations can be airborne (from the sound pressure generated by the loudspeakers) or structure-borne (through the floor or component chassis to which the power transformer is attached). Power transformers vibrate at the power-line frequency of 60Hz or, in some places, 50Hz. The damper is made from a rubberlike material—often Sorbothane—that damps a tube’s vibration and thus improves sound quality. Some dampers are also claimed to extend tube life by acting as heatsinks, radiating heat away from the tube. Tube dampers cost between \$4 and \$25 each. You can see dampers on a tube in [Fig.15-9](#), later in this chapter—they’re the black rings surrounding the tube.

Ground Isolators

Connecting an audio system to a cable-television set-top box may produce a loud hum through the speakers. That's because the cable-TV ground is of a voltage different from the house ground to which your audio system is connected, creating a ground loop. This difference in voltage causes current to flow through the ground, which you hear as hum. The solution is to physically break the connection between the grounds with a small, cylindrical device called a ground isolator or isolation transformer. This inexpensive product is simply a small transformer that passes along the cable-TV signals but breaks the ground connection.

AC Noise Filters

These small devices plug into unused AC outlets in your home, or into unused outlets in an AC conditioner, and filter noise from the AC line ([Fig.15-2](#)).



Fig.15-2 Noise filters plug into unused AC outlets and reduce noise on the AC line. (Courtesy PS Audio and Music Direct)

Replacement AC Wall Outlets

Audiophile-grade AC outlets are available as replacements for the stock outlets in your house. These outlets are described in detail later in this chapter, under “Dedicated AC Lines.”

Audiophile-Grade Fuses

Specialty audiophile fuses are used to replace the stock fuses found inside most audio components. The same size and current rating as stock fuses, these specialty devices reportedly deliver better sound. Their bodies are made of ceramic, for low resonance, and feature internal wiring made from silver for its high conductivity. The endcaps are made of pure silver, with a layer of 24K gold for better signal transfer and protection from corrosion. Many power amplifiers have fuses in the DC supply to the output transistors; theoretically, the system's sound would most benefit from replacing these with fuses of higher quality.

Shorting Plugs

A shorting plug is an RCA plug or XLR connector whose signal and ground conductors have been connected together. Inserted into an unused input of a preamplifier, it shorts the signal conductors to ground to reduce noise.

Bi-Wire Jumpers

Loudspeakers that can be connected with two runs of speaker cable (bi-wiring) are supplied with small plates or wires (jumpers) that connect the two positive terminals together and the two negative terminals together when the speaker is connected by a single run of cable. (The jumpers are removed for bi-wiring.) Replacing a speaker's stock jumpers, which are often made of gold-plated brass, with ones of higher quality can improve the sound. High-performance bi-wire jumpers are often made of solid, high-purity copper plated with rhodium (Fig.15-3).




Fig.15-3 Replacement bi-wire jumpers often sound better than stock brass jumpers. (Courtesy Cardas and Music Direct)

Antistatic Sprays

Supplied in a small aerosol can, antistatic spray applied to equipment and cables reportedly removes built-up charges of static electricity and improves fidelity.

Amplifier Stands

These platforms raise power amplifiers a few inches above the floor, to provide a stable support. The better stands incorporate sophisticated vibration-isolation technology that can significantly improve the sound. Amplifier stands are a much better choice than resting the amplifier on a carpeted floor, because they provide better ventilation as well as vibration control.

Cable Elevators

These devices lift interconnects, loudspeaker cables, and AC cords off the floor. For some reason, getting cables off the floor improves the sound (Fig.15-4). One theory is that floors, particularly floors, particularly when carpeted, store electrostatic charges that adversely affect audio signals flowing through cables. Whatever the reason, cable elevators work, and can also be wedged between cables behind an equipment rack, to separate cables that would otherwise touch. It's particularly important to keep AC cords, digital interconnects, and analog interconnects in close proximity from being laid down parallel to one another. Cables laid in parallel to each other are more likely to couple noise from one cable to the other. That's why, when cables meet, they should be positioned at right angles to each other.

Binding-Post Wrenches

This \$10 accessory should be in the toolkit of every audiophile. It is essentially a double-ended nut driver, with one side for 7/16" posts and one for 1/2" posts. The post wrench makes connecting loudspeaker cables easier, and also makes it easy to apply just the right amount of torque to the binding posts.



Fig.15-4 Cable elevators lift interconnects, speaker cables, and AC cords off the floor for better sound. (Courtesy Shunyata Research)

Vibration Absorbers for Loudspeakers

A few products are available that reduce a speaker's cabinet vibrations and improve sound quality. One device I've found particularly effective is a cloth bag of small but very heavy damping beads—the finished product looks like a bean bag. Placed atop speakers, the bags absorb vibrations, and are especially effective with small, lightweight, stand-mounted speakers. The weight also helps couple the speaker to the stand and the stand to the floor.

Other vibration-control devices address the problem with small discs that reportedly absorb vibration through some internal mechanism, or are made of solid wood. The discs can be placed on components, speakers, listening-room walls, and/or equipment racks. These products tend to be quite expensive, however, and reports of their effectiveness are mixed.

CD Treatments

A wide variety of devices are available that are claimed to improve CD sound quality. These run the gamut from rubber rings that fit around a CD to liquids applied to a CD's surface. Some CD treatments make a worthwhile improvement in the sound; others are of dubious value.

The main category of CD tweaks is the damping device. This can take the form of a rubber ring stretched around the CD's edge; a flat, adhesive ring stuck to the outer perimeter of the CD; or a mat that covers the CD's entire label side. The theory is that the less the disc vibrates, the better the sound. Some hold that a more stable CD results in less activity on the part of the servos in the transport's electromechanical system, which maintain focus, tracking, and speed of rotation. This means less current demand on the power supply.

Be careful about applying anything to a CD that can't be removed. If it degrades the sound—as some may—you're stuck. In addition, some transports simply won't play discs that differ in size or shape or mass from a conventional CD. If you want to experiment with CD rings, use the removable kind. Rings generally cost between \$1 and \$3 each. Try a few—if they don't make an audible improvement, you're out only a few dollars.

CD mats can improve sound quality, but also stress the CD player's or transport's motor system; the motor was designed to spin the mass of a CD, not of a CD and a heavy mat. Some disc-player manufacturers have devised lightweight magnetic clamps that hold the CD firmly in place while adding very little mass.

The second main category of CD treatments comprises products designed to improve the CD system's optical performance. These include liquids applied to the CD surface, and colored paint around the disc edge. How these potions work is a mystery, but some do improve the sound. Again, you must be careful about what you apply to a CD surface: some chemicals can attack the CD's polycarbonate substrate and ruin the disc. Further, wiping the liquid off the disc can create scratches that may cause mistracking on playback.

Finally, CD lens cleaners are available to remove dust and dirt from a player's lens. They are the size of a CD, with tiny brushes that wipe the lens as the disc spins. Lens cleaners cost about \$20. I've never found the need to clean a CD player's lens.

Analog Accessories

Vertical Tracking Force Gauges

An essential accessory for turntable setup, the vertical tracking force gauge (also called a stylus-force gauge) indicates the amount of downward pressure the stylus exerts on the groove. A tracking-force gauge is shown in [Chapter 9](#), on [page 261](#).

Azimuth Meters

Azimuth is the perpendicularity of the stylus to the surface of an LP. If the azimuth is incorrect, one channel will be louder than the other, thus degrading imaging. Azimuth can't be accurately judged by looking at the cantilever in the cartridge body—these are often misaligned at the factory. An azimuth meter is a handheld device with RCA jacks for connecting to the output of your phono stage. While playing a test LP, you adjust the cartridge azimuth until the output levels of the left and right channels are identical. An azimuth meter is the easiest and most accurate way to set azimuth.

Replacement Tonearm Wiring and Phono Cables

Some tonearms allow replacement of the stock wiring inside the arm with an aftermarket cable. This requires removing the cable from inside the tonearm and threading the new cable through the arm. This procedure is best performed by someone experienced in the practice.

Another way to improve the cables in a turntable is to replace the phono cables, the RCA-terminated interconnects that run from the turntable to the phono preamplifier. Phono cables are sometimes terminated with a five-pin JIS (Japanese Industrial Standard) plug at the end that attaches to the tonearm, and RCA plugs at the end that connects to the phono preamp (Fig.15-5).



Fig.15-5 Replacement phono cables terminated with RCA jacks and a five-pin JIS plug (Courtesy AudioQuest)

Record-Cleaning Machines

The best way to keep your valuable LPs sounding good and lasting longer is to clean them with a record-cleaning machine. Unlike a brush, which removes only some of the dirt or redistributes it to other parts of the LP, a record-cleaning machine uses a motor-driven vacuum to extract dirt and grime from the grooves. Although some record brushes can help remove dirt, none approaches the effectiveness of a cleaning machine. A record cleaned with a vacuum cleaner will have far fewer ticks and pops, and a cleaner, more open midrange. Removing the junk in the grooves allows the cartridge to extract more musical information from the record. A clean record will also last longer than a dirty one; the stylus won't constantly be grinding dirt into the groove walls.

The prices of record-cleaning machines vary with the level of automation. Some machines, such as the Record Doctor II by Nitty Gritty, require you to apply the cleaning fluid to the record surface, turn the record over (so the fluid side faces the vacuum slot), and spin the record yourself. Semi-automatic machines spin the record with a motor, but still require manual application of the cleaning fluid and flipping over the LP. Fully automatic machines apply the fluid, spin the record, and even clean both sides at the same time.

The state of this art is the ultrasonic record cleaner. Rather than scrub debris from the grooves with a brush, the ultrasonic cleaner dislodges dirt with very high-frequency soundwaves. Ultrasonic record cleaners are described in detail in [Chapter 9](#).

Cartridge Demagnetizers

Cartridge demagnetizers are often generically called fluxbusters, after the first commercially available demagnetizer, the Sumiko Fluxbuster. These palm-size devices cost less than \$100, and usually have two RCA jacks that you connect to your tonearm cable's RCA plugs. Pushing a button activates a high-frequency signal at the RCA jacks that slowly diminishes in intensity, removing unwanted residual magnetism in the cartridge.

Moving-coil phono cartridges need demagnetizing because the core around which the coils are wound gradually becomes magnetized by the powerful permanent magnets surrounding the coils. Air-cored moving-coil cartridges are not susceptible to accumulated residual magnetism and so don't need demagnetizing.

Moving-magnet phono cartridges build up magnetism on their pole pieces, which also degrades the sound. Note that you must remove the stylus from a moving-magnet cartridge before demagnetizing the cartridge. Otherwise, the demagnetizer will demagnetize the magnets and reduce the cartridge's output level. If you accidentally demagnetize the magnets, get a new stylus: the stylus assembly will also contain new magnets.

Cartridges should be demagnetized after every 50 hours of playing time. Note that demagnetizers designed to bulk-erase tapes should *not* be used to demagnetize cartridges.

Record Brushes

Often made of carbon-fiber bristles, the record brush removes surface dust from LPs, and should be used before playing every LP side. However, the record brush is not a substitute for the periodic cleaning of records with a liquid cleaner.

Replacement Record Sleeves

Paper record sleeves leave small fibers on your records, and plastic sleeves can leach chemicals onto the record's surface. Special rice-paper sleeves produce neither of these problems and will keep your records sounding good for years.

Platter Mats

A platter mat rests atop your turntable platter and is designed to reduce vibration in LPs. Platter mats might be better than no mats on some turntables, but higher-quality turntables are designed so that the least vibration is realized with no mat.

Record Clamps

A record clamp couples the record to the platter so that the record can't vibrate freely. Ideally, the platter acts as a vibration sink, draining vibrations from the record and thus improving the sound. Clamps also help to flatten warped records. Record clamps are described in detail in [Chapter 9](#).

Cartridge Alignment Protractors

A cartridge alignment protractor allows you to correctly set the offset and overhang of a cartridge (these terms are explained in [Chapter 9](#)). A protractor is shown in [Chapter 9](#), on [page 262](#).

Record Flatteners

The record flattener, which looks like a large waffle iron with flat inside surfaces, gently heats your LP while applying even pressure over the entire LP surface, to remove warps from the record.

Stylus Cleaners

Cleaning your stylus before every record is essential not only to good sound, but to preserving record life. A dirty stylus will grind junk into the groove walls that can't be removed, even by a record-cleaning machine. Stylus cleaners come in what look like bottles of nail polish, with a brush attached to the bottle cap. Another type of stylus brush has short, stiff bristles to which a fluid is applied. Whichever brush you use, the stylus should be cleaned with a gentle stroke from the back of the cartridge to the front. A stylus-cleaning kit sells for about \$20.

Antistatic Guns

This device removes the static charge from LPs by creating an opposite charge when the piezoelectric crystal in the gun is squeezed by the trigger.

LP Lubricants

These spray-on treatments reduce the friction between stylus and groove, reportedly resulting in improved sound and less record wear. One application lasts the life of the record.

System Setup Techniques

Equipment Placement

- Provide components with adequate ventilation by maintaining a few inches of air space around each. Avoid stacking components atop one another; overheating will shorten product life. In addition, some digital products are designed to operate at a specific temperature. If an outside heat source, such as a tubed preamplifier, is placed beneath a DAC, the DAC's sound may suffer.
- Maintain adequate distance between the preamplifier and power amplifier—the power amp's large transformer can radiate 60Hz hum. This separation is more crucial for separate phono preamps or for preamplifiers with phono stages.
- Don't position your power amplifier(s) on the floor in front of the loudspeaker plane; that is, the amplifier's front panel should be behind the speakers' baffles. Acoustic reflections from the amplifier can degrade imaging. Also, and especially if the amps are tubed, microphonics can be a problem.
- Install the equipment on a sturdy rack. Vibrations can affect system performance, and particularly that of turntables. (Equipment racks are covered in detail later in this chapter.)
- Avoid placing the equipment rack in a corner where acoustic pressure builds up and could cause components in the rack to vibrate. As explained in the previous chapter, a sound source in a room creates stationary areas of high and low pressure, called standing waves. These high-pressure areas are typically most severe in the room's corners, which is why it's best to position an equipment rack away from the corners.
- If you place an amplifier(s) on a carpeted floor, use spikes, cones, risers, or a stand to raise the amplifier above the carpet, to improve ventilation. The convection cooling essential to keeping the output stage from overheating is compromised if the amplifier's bottom plate sinks into thick carpet (Fig.15-6).



Fig.15-6 Spikes or cones beneath an amplifier or integrated amplifier on a carpeted floor help damp vibrations and provide much-needed ventilation.

Cables and Interconnects

- Because *any* wire degrades the signal passing through it, the less wire you have in your system, the better. Keep analog interconnects and speaker cables as short as possible, but always the same length for the left- and right-channel signals. Note that this doesn't apply to digital interconnects, which sound best when they're between 1 and 1.5 meters long.
- Keep interconnects away from AC power cords. If they must meet, position them at right angles to each other instead of parallel.
- If you must choose between long interconnects and short speaker cables (such as with a preamp in an equipment rack, and power amps on the floor next to the speakers) or short interconnects and long speaker cables (such as with the power amp in the equipment rack), go with shorter speaker cables. There's some controversy about this, but my experience suggests that long speaker cables degrade sound quality more than do long interconnects.
- At the back of an amplifier, where the AC power cord, interconnects, and speaker cables converge, separate the cables as much as possible. If they do meet, cross the cables at right angles instead of letting them parallel each other. Cable elevators designed to lift cables and interconnects off the floor are ideal for separating cables behind an amp (Fig.15-7).




Fig.15-7 Separate AC cords, interconnects, and loudspeaker cables where they meet behind an amplifier. Use the same technique for interconnects and power cords behind an equipment rack.

- Lift cables and interconnects off the floor with cable elevators.
- If you have an excess length of cable or interconnect, don't wind it into a neat loop behind the speaker or equipment rack. This will make the cable more inductive and change its characteristics. Instead, drape the cable so that it crosses other cables at right angles.
- Periodically disconnect and clean all interconnects and speaker cables (Fig.15-8). Oxidation builds up on jacks and plugs, degrading the audio signal passing through them. Use a contact cleaner and contact enhancer (described earlier). Begin by cleaning all the RCA and XLR jacks on your equipment and cables. Treat the outer ring, as well as the unseen junction inside the jack that contacts the RCA plug's center pin. Pay special attention to phono jacks and plugs; connectors that carry the tiny phono signals from the cartridge to your preamplifier are the most sensitive to degradation from dirt and oxidation. Next, clean the binding posts on speakers and power amplifiers, as well as the spade lugs on speaker cables. Be sure not to touch any cleaned surface with your fingers—body oils will trap contaminants. In fact, it's a good idea to wear gloves when cleaning contacts. You may also consider cleaning the pins of phono cartridges, and the tonearm's sleeve connectors that attach to them. After cleaning, apply a

contact enhancer to pins, jacks, binding posts, and spade lugs. Periodic cleaning and enhancing of all electrical contacts in a system can keep it sounding its best—say, once every two months. Make sure, however, that the contact cleaner doesn't leave a residue that could attract dust and end up degrading the sound.



Fig.15-8 Periodically clean spade lugs, binding posts, RCA plugs, RCA jacks, XLR jacks, XLR plugs, tube pins, tube sockets, and other connections with contact cleaner.

- When connecting and disconnecting RCA plugs, always grip the plug, never the cable. When disconnecting male XLR plugs, remember to first push the tab on the female jack.
- Ensure tight connection of all RCA plugs, particularly the spade lugs of power amplifiers and loudspeakers. Make sure you get lots of contact surface area between spade lug and post, then tighten down the binding post's nut.
- Avoid sharp bends in cables and interconnects—kinks and bends can deform the cable's geometry (the physical relationship between the conductors and surrounding insulation), and thus degrade their performance.
- When connecting speaker cables, a high contact pressure is more important than a large contact area. A tight connection—enough to slightly distort a thin

spade lug—is ideal. Be careful not to overtighten the binding-post nut, which could cause the entire post to turn inside the speaker, possibly breaking the connection inside the speaker between the binding post and the speaker’s internal wiring.

- Once cables and interconnects are in place, don’t disturb them. Cables and interconnects seem to “settle” and sound better over time, a process that must begin anew if the cables are moved. Cables and interconnects should be moved only for cleaning.
- If you have the choice, opt for spade lugs over banana plugs—spades provide better contact with the binding post.

Tubed Equipment

- Vacuum tubes are microphonic—they can act as microphones, creating an electrical signal when physically vibrated. Prevent tube microphonics by encircling each tube with specially made damper rings (Fig.15-9).
- Keep the tubes biased according to the amplifier manufacturer’s specifications.
- When installing or changing tubes, wear cotton gloves to avoid depositing body oils on the tubes’ glass envelopes (Fig.15-9). The oil on your fingers can damage the glass.



Fig.15-9 Wear gloves when installing or replacing vacuum tubes. Note the tube dampers, the two black rings surrounding the tube.

- Clean the contact points between the tubes and the socket. Disconnect the power to the component, remove the tubes, and clean the contacts on the tube sockets, using an audio contact cleaner and a Q-tip. Then clean the tube’s pins and treat them with contact enhancer. Reinsert the tubes.
- Tubed amplifiers often offer different transformer taps that allow you to select the tap that best matches your speakers’ impedance. The taps are often marked “8 ohms” and “4 ohms,” to suggest the optimal taps for 8-ohm and 4-ohm speakers.

Technically, the taps allow you to select the transformer's turns ratio, a term that describes the ratio between the number of turns of wire in the output transformer's primary and secondary windings. A higher turns ratio produces a higher output voltage, but less current. A lower turns ratio produces a lower voltage, but higher current. Maximum power transfer occurs when the correct turns ratio for a given speaker's impedance is selected. However, the tap that provides the greatest power transfer might not produce the best sound. The surest way to decide which tap to use is to ignore the speakers' nominal-impedance specification and just listen to all the taps.

- Tubes wear out—their performance degrades over time, this decline manifested as a lack of treble sparkle, duller reproduction of transient information, and reduced power output. The output tubes in power amplifiers wear out much faster than low-level tubes in preamplifiers and the input and driver stages in power amplifiers.

LP Playback

- When playing LPs, turn off your digital components (CD players, DACs, music servers). Digital circuits radiate noise that can contaminate the audio signal.
- Periodically give your stylus a deep cleaning: Hold a matchbook below the stylus, and very gently scrape the stylus over the matchbook's striking area by pulling the matchbook toward you, from the back of the cartridge to the front—that is, in the same direction as the groove moves under the stylus. This abrasive action removes vinyl that has accumulated on the stylus.
- Clean LPs—even new ones—with a liquid cleaning solution. LPs are pressed with a mold-release compound that makes them easier to remove from the stamper; this should be removed before they're played. Vacuum or ultrasonic record cleaners are optimal, but a brush and liquid solution will also work.

Digital Playback

- When connecting your source components to your preamplifier, keep separate from each other digital data interconnects (S/PDIF, AES/EBU, USB) and analog interconnects. The very high frequencies carried by digital cables can radiate noise and pollute analog signals.
- If your CD player or DAC has a digital-domain volume control, run it with 3dB of attenuation rather than at maximum level. Although attenuating the volume with a digital-domain volume control theoretically reduces resolution, most digital filters sound better when not processing full-scale signals. The very slight reduction in resolution is more than made up for by the greater sense of ease when a few dB of attenuation are applied.
- Warm up your digital components for at least two hours before listening. Longer is better—it can take a full day for a DAC to sound its best. I leave my DAC powered up continuously.

Loudspeakers

- Periodically check the tightness of the driver mounting bolts (Fig.15-10). These can come loose over time, degrading performance. They also can loosen during shipping, so check their tightness when you first unbox new speakers. If the driver's frame isn't tightly coupled to the baffle, the frame will vibrate and produce a sound of its own. But don't overtighten the bolts.



Fig.15-10 Periodically tighten the bolts holding the loudspeaker drivers to the baffle.

- Ensure good contact between speaker cables and binding posts. Get the spade lug fully on the binding post before tightening it with a binding-post wrench.
- Be sure that speakers don't rock on their spikes. All four (or three) of each speaker's spikes should be fully coupled to the floor.
- If you don't want to damage hardwood or tile floors with spikes, buy a set of round discs with small indentations in the center that are made specifically for this purpose and insert the discs between the spike and the floor. If you use coins rather than specially made discs, use quarters for their greater thickness.
- When placing small speakers on stands, couple the speaker to the stand with four small balls of Bostik Blu-Tack.

- Using a bubble level, check that each speaker has the same degree of tilt-back (also called rake angle). This is particularly important with speakers that have a first-order crossover.
- Bi-wiring speakers with two different types of cable requires great care. Every cable has its own capacitance and inductance, which interact with the capacitance and inductance of the speaker's crossover network. Using cables that have different characteristics relative to each other can change the crossover's performance.
- To break in speakers more quickly, play them at a somewhat loud level and note the volume setting on the preamplifier. Turn the system off entirely. Face the speakers directly toward each other, about 1" apart. Reverse the red and black leads on only one speaker, and throw a blanket over both speakers. Power up the system. Play a test CD or digital file of wideband pink noise at a level of -10dB, and put that track on Repeat. This technique will maximally break in the speakers with minimal audible noise. With the speakers facing each other and driven out of phase, their outputs will largely cancel, thus reducing the sound level, although the drivers will still be worked hard; the blanket will damp any remaining sound. Remember to determine the volume level before wiring the speakers out of phase; when they're out of phase, it will be impossible to judge how loudly the speakers are playing, making it easy to overdrive and damage them.
- Remove speaker grilles. Although some manufacturers claim that their grilles don't degrade fidelity, virtually all of them do by creating diffraction, which is never desirable.

AC Power

- Don't use light dimmers or fluorescent lighting in the listening room. Dimmers work by chopping up the 60Hz AC sinewave, which creates lots of noise in the AC line, and this noise gets into your components through their power cords. If the listening room has dimmers, turn them all the way off.
- When evaluating power cords, replace every cord in your system with the cord under evaluation, not just on one component. The degradation imposed by all the other cords can easily swamp the benefits of a single superior cord.
- Replace the stock AC wall sockets with audio-grade models (Fig.15-15). Use contact enhancer on the blades of AC plugs.
- Have an electrician tune up your home's AC power: 1) Tighten the screws holding the circuit breakers to the panel and to the bus bars. 2) Replace the breakers that supply your audio system; they probably have carbon deposits, which can add noise to the line. 3) While you're at it, have the electrician replace 15-amp breakers with 20-amp units. 4) Clean and reseal the home's main ground wire to the grounding rod or copper water pipe.

General

- Periodically, with the system turned off, turn your preamplifier's volume control up and down several times. This action cleans the volume control's wiper. This applies only to preamplifiers that have a continuously variable resistive volume control, and not to models with a stepped attenuator or those with digitally controlled volume adjustment. If the volume is adjusted in steps, or if the preamplifier's display offers a numerical readout of the volume level, the preamp doesn't have a continuously variable resistive volume control.
- Warm up your system before a listening session. An hour is a good rule of thumb, although different components have different warm-up times. Large solid-state amplifiers seem to take as much as two hours to sound their best, and digital gear needs at least 12 hours (preferably 24) to reach optimal performance.
- When listening, turn off all fluorescent lights, including the newer compact fluorescent bulbs. All fluorescents employ a small switching circuit that radiates noise.

Equipment Racks and Isolation Devices

Equipment Racks

The most important accessory to any audio system is the equipment rack. A good rack presents your components in an attractive way, makes your system functional and easy to use, and, most important, helps you get the best sound possible from your system. If your system includes a turntable, a high-quality rack is absolutely essential to realizing the LP front end's musical potential.

Apart from providing convenient housing for your equipment, the rack's main job is to isolate the equipment from vibrations. There is no question that vibrations degrade the sonic performances of preamplifiers, DACs, disc players, and especially turntables. These vibrations are generated by transformers in power supplies, motors in turntables and CD transports, and from acoustic energy impinging on the electronics.

Vibration affects electronic products through a phenomenon called *microphony*. A microphonic component emits a small electrical signal when vibrated; in other words, it acts as a microphone, converting mechanical energy into electrical energy. This electrical energy pollutes the audio signal, resulting in degraded musical performance. Tubed products are particularly susceptible to microphony. In fact, you can shout into some tubed preamplifiers and hear your voice reproduced by the loudspeakers—the tubes are converting the acoustic energy of your voice into an electrical signal. When the tubed component is vibrated by the sound from the speakers, or by vibrations from transformers or motors, noise is added to the music.

The sound qualities of digital products, too, can be degraded by vibration, but through a different mechanism. The crystal oscillators in many DACs change their frequencies slightly when vibrated. Although this shift in frequency is very slight, even small variations in timing can wreak havoc in a digital audio system. Some manufacturers of digital products go to extreme lengths to ensure accurate and stable crystal frequencies, and build some form of vibration isolation around the oscillators.

Most sensitive to vibration is the LP playback system—vibrational energy is transmitted to the tonearm and cartridge, adding sonic garbage to the tiny signals recovered from the record groove. These vibrations come from the turntable motor, from resonances in the tonearm, and from acoustic energy impinging on

the turntable, arm, and cartridge. Isolating your system from vibration is essential to realizing the best musical performance.

A good equipment rack fights vibration with rigidity, mass, and careful design, including technologies that dissipate vibrations. The massive, inert structure of a high-quality rack is much less likely to vibrate in sympathy with the sound pressure generated by the speakers. Moreover, the equipment rack can absorb, or damp, the vibrations created by power transformers and motors. The shelves of many equipment racks include built-in vibration-damping mechanisms.

Avoid racks with large, unsupported shelves, flimsy construction, low mass, and/or poor vibrational damping. And don't even consider the generic "stereo stands" sold in furniture and department stores—they aren't designed to combat vibration. Get a rack specially designed for high-performance audio systems. When choosing an equipment rack, look for a thick steel frame (14-gauge steel is excellent), welded construction, the ability of the leg posts to be filled with sand, shelves made of a vibration-damping material (e.g., medium-density fiberboard, aka MDF), and floor spikes. This description is of high-quality audio racks, but some very high-end racks feature much more sophisticated design and construction. Virtually all such racks are supplied with spikes pointing down from the rack legs to the floor. A spike acts as a mechanical diode, allowing vibrational energy to flow in only one direction, from the rack to the floor. For a spiked rack to be effective the floor must be sturdy, and flat so that the rack doesn't rock. The heights of most spikes are adjustable; you can level the rack and get good contact between the floor and all four (or three) spikes. Rack spikes are usually much more heavily built than speaker spikes, with more rounded points. Note, however, that both types of spike will damage wood floors. Some racks are supplied with small plates to hold the spikes, to prevent such damage.

Not every good rack will have all of these features, but most will have at least some. The ability to fill the rack's posts with sand is especially important, unless the rack has built-in mass and damping. Sand is an excellent vibration-absorbing material, and it increases a rack's mass. It's also cheap.

Since the previous edition of this book, a number of manufacturers have entered the market with extremely high-tech racks designed using aerospace engineering. These racks are elaborately designed using the same computer modeling programs used to study vibrational analysis in other fields. These products, while expensive, have taken isolation from vibrations to unprecedented levels of sonic performance (Fig.15-11).



Fig.15-11 A high-quality equipment rack reduces vibration in your components. (Courtesy Critical Mass Systems)

The preceding discussion of high-mass racks, floor spikes, and intimate coupling of components to racks and floors is only one school of thought. Another theory holds that audio components should be decoupled from their support surfaces and the floor with lightweight racks and compliant feet. Vibrations in the rack and components are thus damped rather than transmitted to the floor. This is, however, the minority view.

Before installing a heavy rack—particularly one to be filled with sand—make sure the floor structure can support the weight. Suspended wood floors may need additional bracing, which will also make the floor less likely to vibrate. A concrete slab is the ideal foundation for an equipment rack; you don't have to worry about the weight, and it provides a solid surface for the rack. Before filling a rack's posts with sand, be sure you've set it up exactly where you want it.

Vibration-Control Accessories: Spikes, Feet, and Cones

Vibration-isolation accessories and specialized audio equipment racks are designed to do the same thing: reduce vibration in audio components. Isolation devices can be used on their own, or in conjunction with an equipment rack.

The field of vibration-control accessories for audio systems was established with a product called Tiptoes. Invented in 1983 by Steve McCormack, founder of The Mod Squad, Tiptoes are machined aluminum cones that end in sharp points. When used point-side down to support pieces of audio equipment, they act as mechanical diodes to remove mechanical energy from the components. When introduced, Tiptoes met with scorn—until their effects were heard. Tiptoes spawned an entire industry of cones, spikes, and similar devices of all manner of construction, made from a wide range of sometimes exotic materials. Some sophisticated isolation devices employ tiny mechanisms inside them that convert vibrations into minute amounts of heat (Fig.15-12).




Fig.15-12 Spikes, cones, and compliant feet. (Courtesy AudioPoints, Stillpoints, and Vibrapod/Music Direct)

A different approach to vibration control is to support the component with compliant feet, or mount the component on a compliant platform. These devices are made of a rubber-like material (usually Sorbothane) chosen for its excellent damping characteristics. Compliant feet isolate the audio product from external vibration by absorbing the energy before it can reach the product. These devices are also excellent for stacking one component atop another. The feet provide airspace for ventilation, and the additional space reduces potential magnetic interaction between components.

Since the last edition of this book (2010), I've had extensive experience with state-of-the-art vibration-control devices and equipment racks. I can tell you that,

as important as I thought vibration control was, hearing the cumulative effect of elaborate vibration control in my system has led me to believe that vibration is a significant source of sonic degradation. Every upgrade I've made in my system's ability to minimize vibrations—**isolation** devices, equipment racks, amplifier stands—has resulted in an important increase in transparency, resolution, coherence, dynamics, bass performance, and soundstage dimensionality.

AC Power Conditioners and AC Cords

Over the past ten years, the audio world has learned much about the role AC power plays in the quality of an audio system's sound. We had AC power conditioners before that time, of course, but many of those early conditioners—and, indeed, many of today's models—take a hit-or-miss approach. Their designers attempted to address the problems of AC power without fully understanding the phenomena involved, and consequently often created products that did as much harm as good (many compress music's dynamic range). But some of today's power conditioners and cords are far more sophisticated and effective than those of a decade ago, and are based on a firmer understanding of how AC power interacts with the audio signal to affect a system's sound. Power-conditioner design has, in some cases, been transformed from a trial-and-error experiment into a serious engineering exercise based on genuine scientific research. These conditioners are vastly better than those of the previous generation at making possible significant improvements in a system's sound quality.

An AC power conditioner plugs into the wall outlet and provides multiple AC outlets for plugging in your audio equipment. But before we talk about what a power conditioner does, let's look at the AC power line itself, and its relationship to an audio system. The AC power from the wall outlet is a 60Hz, 120V sinewave that powers the audio system (see [Appendix B](#) for more details about the power line). All of your components are connected to each other via the power line. In fact, your audio components are connected to every other electrical device in your house, and to every home and factory also using the power grid.

Any appliance plugged into the wall generates noise on the AC line. This noise then appears on every other component plugged into your home's AC outlets. The amount of this noise, called electromagnetic interference, or EMI, varies with the appliance. Light dimmers, refrigerators, vacuum cleaners, and electrical power tools are major sources of powerline noise because the fibers in the motors' brushes are continually making and breaking contact, putting back on the line a small transient spike with every contact. The line is also polluted by AM radio stations; power lines act as antennas, superimposing the AM signal on the 60Hz line. This noise on the AC line then gets into your audio components, where, among other problems, it adds a layer of grunge to instrumental timbres, hardens the treble, reduces low-level resolution, and obscures very fine spatial

cues, which diminishes the soundstage size and dimensionality of instruments within the soundstage.

Because the designers of early power conditioners assumed that this noise on the AC line was the most significant power-line problem facing audio systems, their products filtered incoming noise from the line. Most improved the sound quality to some degree, but these designers had ignored another significant source of noise on the AC line: the noise generated by other audio components in the system that were plugged into the conditioner. The early conditioners isolated the audio system from outside noise, but failed to address the noise put on the line by the audio components themselves.

How do your audio components create noise on the AC line? It's natural to assume that, when your system is running, AC is being continuously pulled from your wall. In this view, AC power is a low-frequency (50 or 60Hz) event, and this is indeed the case with fans, motors, and other constant-draw devices. But an audio system makes very different demands on the AC power system. The capacitors in an audio component's power supply charge—that is, draw current from the wall—only in very short bursts, at the tops and bottoms of the power line's 60Hz AC sinewave. This means that, 120 times per second, the AC power system delivers a brief but high-level burst of current that recharges the audio component's capacitors. These bursts of current put noise back into the power line. Specifically, at each burst of current, diodes in the bridge rectifiers inside each audio component turn on and off, radiating noise. The effect is particularly pronounced in power amplifiers, whose current requirements are much higher than those of other components. Another source of noise is any component with digital circuitry, which includes not only the obvious products such as DACs, but also the increasing number of preamplifiers controlled by software. This problem is greatly exacerbated by the trend toward incorporating into our audio systems computer-based music servers, which are horrendously noisy. Connecting a computer to a DAC via the USB cable couples the computer's very noisy ground to the DAC, and then on to your other components.

It's not enough to keep noise from the AC line from getting into your audio system; just as important is to isolate each audio component from the others. Today's best power conditioners reduce this noise coupling between components with filtering in the power conditioner, and even within the power cords themselves. [Fig.15-13](#) shows how an AC line conditioner prevents noise in one component (in this case, the CDP or compact-disc player) from getting into others. [Fig.15-14](#) shows the rear panel of a high-quality AC conditioner.

CCI — CAPACITIVE FILTERING

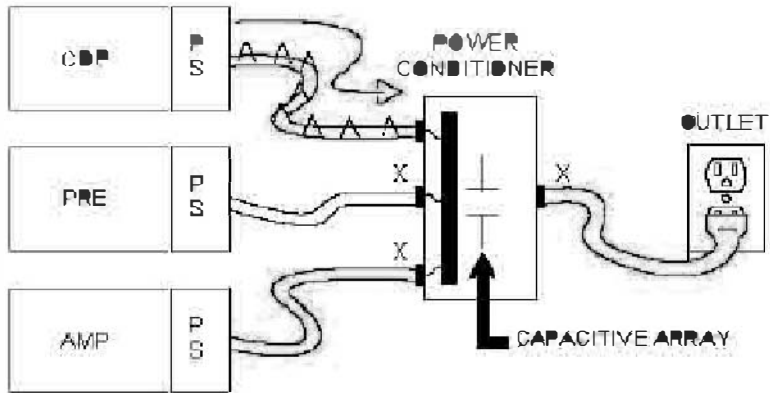


FIG. 5

Fig.15-13 How an AC conditioner isolates each component from noise (Courtesy Shunyata Research)



Fig.15-14 A high-quality AC power conditioner (Courtesy Shunyata Research)

Isolating the system from noise is only half of the AC conditioner's job. The AC conditioner and power cord should also provide a low-impedance path for the current. Consequently, some AC conditioners are built with large, solid-copper bus bars, custom AC sockets that provide more contact area and greater contact pressure with the plug, and other techniques to provide maximum current delivery. These parts are often cryogenically treated by subjecting them to very low temperatures, then slowly bringing them up to room temperature. Cryogenic treatment alters a metal's molecular structure and, for some reason, results in better sound quality. (I've heard comparisons of cryogenically treated and non-treated interconnects, and have been surprised by the difference.) Earlier-generation power conditioners, and many of today's models, often compressed music's dynamic range because they restricted current flow. Today's best power conditioners are often designed in conjunction with AC cords as complete systems for minimizing noise and maximizing power transfer.

A mistake commonly made in the evaluation of AC cords is to install only a single cord in a system. The new AC cord could be spectacularly great, but you might not know it if all of the other AC cords in your system are compromised. That's why it's essential to replace all AC cords at once before reaching any judgments about an AC cord's quality. When choosing a line conditioner, make sure its power capability exceeds the total power consumption of all the components you'll be plugging into it. Choose a conditioner with enough outlets for your present and your anticipated needs. And, as with all accessories, try the power conditioner in your system before you buy. Expect to pay a minimum of \$250 for a conditioner with just a few outlets, or several thousand dollars for a state-of-the-art conditioner with many outlets.

A radical approach to AC line conditioning is to create an entirely new 60Hz sinewave by amplifying the output of a small 60Hz oscillator. This sort of power conditioner is essentially a power amplifier that operates at only one frequency: 60Hz. As you can imagine, this type of conditioner's output power is limited, making it best suited to frontend components rather than power amplifiers. These products are large, heavy, run hot, and—unlike a power amplifier, whose output depends on the audio signal's amplitude—consume just as much power when your audio system is idle as when it is playing music. But while this technique creates perfectly clean power, it doesn't address noise put onto the local (i.e., within your audio system) AC line by your components themselves.

A power-line conditioner can't make poor audio components sound good; it can only provide the optimal AC environment for those components, so that they can realize their full potentials. The sonic benefits of a good line conditioner

include a quieter background, with less low-level grunge and noise. The music seems to emerge from a perfectly quiet, entirely “black” space from dead silence, rather than from a grayish background. The treble often becomes sweeter, less grainy, more extended. Soundstaging often improves, with greater transparency, tighter image focus, and newfound depth. Midrange textures become more liquid, and the sound has an ease and musicality not heard without the conditioner.

AC Power Cords

At first glance, it may seem that the several short pieces of wire carrying AC to the components in your audio system—the AC power cords—couldn't affect the sound of the system. After all, the AC cord is just another piece of wire at the end of a miles-long transmission system; how could it make any audible difference?

From the perspective of your audio components, the AC cord isn't the last few feet of a transmission system, but the *first* few feet. The AC cord is directly connected to the primary windings in the power transformer of your audio component, and because all the components are powered by the same AC line, noise from one component easily travels to all the other components via the AC line. Good AC power-cord design minimizes the noise coupling between components described earlier, and ensures maximum instantaneous current delivery. Replacing all the stock AC cords with entry-level cords is probably the single most cost-effective sonic upgrade you can make to your system. Generic AC cords that are supplied with audio components are anathema because they limit instantaneous current delivery and have no provision for reducing noise transmission between components.

Dedicated AC Lines

The ultimate upgrade of your power is dedicated AC lines—lines that feed your audio system but are independent of the lines supplying power to the rest of your house. Before going any further, I must unequivocally state that installing dedicated lines must be performed only by a licensed electrician. Do not attempt to modify your home's AC system yourself.

How can dedicated AC lines improve an audio system's sound quality? To answer that question, let's look at how an audio system interacts with conventional house wiring. In most homes, the lengths of wire between the AC outlets and the breaker panel vary with the outlets' distances from the panel. If you plug a preamplifier into one outlet and a power amplifier into a different outlet, each component will "see" a different length of wire to ground, and thus a different resistance between the component and the ground. This disparity in ground resistance can result in a very slight voltage difference between the two grounds, and create current flow called a ground loop. This ground-current flow is minuscule, and doesn't affect the performance of lights, vacuum cleaners, and toasters, but it can degrade an audio system's sound quality by raising the system's noise floor and obscuring low-level detail. In addition, ground loops cause the music to emerge from a background that is less than perfectly silent, and thus diminish the sense of realism. You don't hear the noise overtly, as audible noise, but rather as a "grayness" to the background. In extreme cases, this ground-loop current is heard as a 60Hz hum through your speakers. Some audio components offer a ground-lift switch that disconnects the component from the ground, to avoid this problem. If a component has a ground lift switch, listen to the system with the switch in both positions and leave it in the position that produces the quietest background.

Another shortcoming of conventional AC wiring is the house electrical system's inability to deliver adequate current to the audio system. The amount of current the AC system can deliver is determined by the capacity of the circuit breakers in the breaker panel and the thickness of the wire between the breaker panel and the AC outlets. These specifications were determined decades ago, and specified in the electrical code, by considering the current draw of an average number of standard household appliances, then specifying a wire thickness and breaker rating that could safely deliver enough power to those appliances.

But an audio system makes very different demands on an AC power system than does a refrigerator or a fan. Nearly all household appliances draw an

unvarying stream of current from the wall—think of a light bulb. Conversely, an audio component pulls power from the wall only in very short bursts, during the peaks of the power line's 60Hz sine wave. For a power amplifier, these bursts of current flow are quite intense, as the inrushing current recharges the amplifier's large reservoir (filter) capacitors. In fact, that's why those capacitors are needed—they supply the amplifier with power during the relatively long periods between the brief bursts of current draw from the AC line.

Conventional AC house wiring wasn't designed to supply the high peak current demands of an audio power amplifier. Although the bursts are too brief to heat the wire and create an unsafe condition, the wire and the breaker nonetheless restrict current flow to the amplifier. This problem is compounded by cheap-as-possible AC outlets that provide little contact area between outlet and plug.

The solution to all these problems is to replace the AC supplying your audio system with dedicated AC lines. Typically, two or more lines are installed: one line for the frontend components and one for the power amplifiers. My own audio AC system comprises four dedicated lines: one to power the analog source components, one to power the digital sources and preamplifier, and one for each monoblock power amplifier.

I emphasized at the outset that any work on a home's AC system must be performed by a licensed electrician. Unfortunately, electricians are largely unaware of the unique demands an audio system places on the AC power system; they assume that any job done "to code" is good enough, and that the techniques for improving an AC power system I'm about to describe are unnecessary overkill. Half the battle in improving your AC power is getting your electrician to understand why these procedures are vital to perfectionist audio, and to follow your instructions.

Here's what I learned about dedicated lines in my research and consultation with industry experts on the topic before having the lines installed in my home. A little later, I'll describe the sonic differences the dedicated lines made.

The wires between the breaker panel and the outlets must be of the same type, size, and length, and be continuous (i.e., no splices). If one outlet is farther from the panel than another outlet, the closer outlet must necessarily have some excess wiring in the wall. Identical wire lengths ensure that each outlet's ground path is identical and that no ground loops are created. The wire must be copper, not aluminum, and be 10 AWG or thicker. (A wire's thickness is described by its

American Wire Gauge number; the lower the AWG number, the thicker the wire.)

If the main breaker panel is more than 35' from the outlets, have the electrician install a subpanel near the audio system. The subpanel should be powered by a single phase from the main panel. The wiring between the main panel and the subpanel should be 8 or 6AWG. The longer the run of wire, the thicker the wire needs to be. Even if your electrician tells you that 10AWG is sufficient, specify the thicker wire.

Replace any breaker supplying your audio system, in either the main panel or subpanel, with silver-tungsten (not copper) contacts. Copper-contact breakers have higher resistance and will corrode. Use silver paste on the breakers' screw terminals, and on the contacts in the AC outlets that attach to the wire. This paste reduces the resistance at the contact points, and increases the AC system's ability to deliver current. McMaster #1219K57 silver paste is ideal. The electrician should ensure that all contacts between breakers and wire are clean, tight, and treated with silver paste.

Choose AC outlets specifically designed for audio systems. The cutaway photos in Fig.15-15 show a conventional AC outlet (top) and one designed for audio (bottom). Although the two outlets appear superficially similar, the high-end outlet has a more robust method of attaching the wire to the outlet, greater surface-contact area between the outlet and the inserted plug, thicker conductors, and better metallurgy that has been designed for maximum conduction. The high-end outlet also provides a tighter and more secure fit of the AC cord's blades in the socket. It's a good idea to ask the electrician to install two duplex outlets on each dedicated line so that you can plug multiple components into the line. Finally, every component in your audio or audio/video system must be powered from the dedicated lines.




Fig.15-15 Comparison of a standard grade AC outlet (top) and one designed for audio performance (bottom) (Courtesy Shunyata Research)

How a dedicated AC system is grounded will have a large effect on its performance. If the dedicated lines' ground is tied to the house ground (usually, the incoming copper water pipe), have the electrician clean the pipe, the home's ground wire, and the clamp that holds them together. That green film of corrosion that appears on copper pipes acts as an insulator to degrade the ground's performance—it must be sanded off before the ground wires are attached. Have the electrician apply silver paste to all contact surfaces before

clamping the ground to the copper pipe. The idea is to create a very low-resistance path to ground.

The ultimate ground for your audio system is a ground rod driven into the earth that is isolated from the home's ground. The ground rod should be in dirt that will remain moist, such as an irrigated garden. I've seen unbelievably elaborate grounding systems in which the ground rod is attached to a grid buried in the earth, and is surrounded by a wet electrolytic slurry for better conductivity between the ground wire and the earth. The 2' x 2' grid buried in the earth provides a lower resistance path to ground than a single rod.

Dedicated lines don't obviate the need for AC power conditioners and cords. I also installed an elaborate state-of-the-art AC power-conditioning system and AC cords, giving me the ability to independently evaluate the effects of the dedicated lines, the power conditioners, and power cords. I left my house's stock AC outlets and wiring in place so that I could perform comparisons with the new dedicated lines. The stock-wired outlets and the new dedicated outlets were just 16" apart, making it easy to move the AC cables between the stock and dedicated lines for listening comparisons.

The dedicated lines rendered a wholesale improvement in my system's sound. The bottom end was deeper, tighter, and more authoritative. Kick drums had greater dynamic impact and punch. This quality alone would have been worth the price (about \$2400 in my installation), but the dedicated lines also lowered the noise floor, revealing more low-level detail, including spatial cues. Consequently, the soundstage became larger, with a greater sense of three-dimensional layering, and with more precise image specificity. The dedicated AC lines increased the system's resolution, allowing me to more easily hear differences between components under review. Overall, I'd say that installing dedicated AC lines was the best audio investment I've ever made.

A Final Note

The accessories and techniques described in this chapter are designed to improve the sound of your system so that you can more easily forget about the equipment and enjoy the music. When listening for pleasure (as opposed to critical listening to evaluate an accessory or technique), forget everything you've read in this chapter—the last things that should be on your mind are whether the accessory you've just installed is an improvement, which accessory you're going to try next, or whether the system would sound better with more tweaking. It would be a supreme irony if the very products and techniques created to take you a step closer to the music ended up distracting you from the real reason any of us owns a highperformance audio system: to enjoy the music.

16 Specifications and Measurements

An audio component's technical specifications and measured performance can tell us much about the component. Understanding what the specifications mean, and interpreting laboratory-test results are useful to the audiophile, though not essential. Specifications and measurements have their place in audio, but they are not the beginning and end of the story. Rather, they are just one facet of a multidimensional puzzle.

What specifications and measurements cannot tell us is how the product communicates musical expression. To discover that, we must listen. Although understanding this chapter is not essential to enjoying high-quality music reproduction in your home, I've included an explanation of specifications and measurements for each component category for those readers interested, and to make this book complete.

Preamplifiers

An important preamplifier specification is a line-stage's *gain*, or the amount of amplification the preamp provides. Gain is expressed either in decibels (dB), or as a number representing the ratio between the input and output voltages.

The preamp's gain is important to consider in relation to the output voltage of source components. If you have a very-high-output CD player or DAC feeding a high-gain preamplifier, you'll barely crack open the volume control before you get loud listening levels. This has several disadvantages. First, it makes setting a "just right" volume difficult: The volume control is rendered so insensitive by such high outputs that even small adjustments make large differences in listening level. Second, volume controls are often less well matched between channels at the low end of their range. This means that one channel becomes louder than the other when the volume control is turned down. The effect is the same as if you'd turned a balance control off-center. This difference can be as much as 3dB—enough to shift the image to one side. A DAC putting out, say, 5V RMS at full scale and driving a line-stage preamp with 25dB of gain would produce this problem.

High gain can be an advantage when using a line-stage preamplifier with an external phono stage. Many phono stages have just enough gain of their own to feed a line-stage, but no more. The line-stage therefore requires a fair amount of gain.

Line-stage gain varies between 5dB and 25dB. The 5dB figure is very low—perhaps too low for some outboard phono stages, and even for tuners and tape decks. This will be especially true if the power amplifier has low gain or the system has insensitive loudspeakers, and could require using the volume control near the maximum setting. For most CD sources, about 10dB of gain is ideal.

As described in [Chapter 4](#), phono-stage gain should be matched to your phono cartridge. Moving-magnet cartridges need about 35dB of gain, while moving-coil types require up to 65dB. If the phono preamp has less than adequate gain, a line-stage preamp with high gain can further amplify the phono signal for adequate listening levels.

A preamplifier's gain can be expressed indirectly as *input sensitivity*. This is the voltage required to produce some reference level (usually 0.5V) at the preamplifier output with the volume turned all the way up. A preamplifier with

an input sensitivity of 0.05V has a gain of 10 (which can also be expressed as 20dB, because 20dB represents a 10-to-1 voltage ratio).

A preamplifier's *input impedance* specification describes the electrical resistance to current flow "seen" by the source component (such as a CD player) driving the preamp. Most preamps have input impedances between 10,000 and 50,000 ohms (10–50k ohms). The preamp's input impedance is the *load* the source component drives. As described earlier, the higher the input impedance, the easier it is for the source to drive. As the input impedance increases, however, current flow through the interconnects and preamplifier input stage decreases. Some designers believe moderate rather than very high input impedance results in the best sonic performance. If the preamplifier's input impedance is too low, it can *load down* the source component, resulting in poor technical and sonic performance.

A preamplifier is also specified according to its *output impedance*. Output impedance is best thought of as a resistor inside the preamplifier that couples its output stage to the rear-panel jacks. Output impedance describes the preamplifier's ability to deliver current into a load (the power amplifier). A preamplifier's output impedance is much lower than its input impedance; where input impedance may be 50k ohms or more, output impedance may be less than a thousandth of that value.

By having a high input impedance and a low output impedance, the preamplifier acts as a buffer between the source components and the power amplifier. The sources drive a high input impedance—which is very easy for them—and the preamp takes on the burden of driving the interconnects and power amplifier. The preamplifier's low output impedance makes it ideal for this job. (The advantages of low output impedance are discussed in more detail in [Appendix B](#).)

A preamplifier's *input overload* specification refers to the maximum input voltage the preamp can accept without distorting (with distortion defined as 1% total harmonic distortion, or THD). The higher the input overload figure, the better. A full-function preamplifier will have two input-overload specifications: one for the line-level inputs (typically more than 10V RMS), and one for the phono inputs (typically several hundred millivolts). *Maximum output level* is the maximum voltage the preamplifier can swing from its output stage. Maximum output level is often as much as 50V—far higher than any audio signal. Because most power amplifiers have an *input sensitivity* (the input voltage required to produce maximum power output) of under 1.5V, virtually any preamp has enough output voltage to drive any power amplifier.

Signal-to-noise ratio, or *S/N ratio*, is a measurement of how quiet the preamplifier is. It expresses in decibels the ratio between a signal level of 0.5V output and the residual noise floor. The higher the signal/noise ratio, the quieter the preamp. Because the ear isn't equally sensitive to noise at all frequencies, a *weighting curve* is sometimes applied to the noise spectrum to more closely approximate the noise's audibility. (Weighting curves are described in more detail in [Appendix A](#).) A signal/noise ratio of 90dB unweighted (no weighting curve applied) is good; higher signal/noise ratios are common when "A"-weighted (some of the noise is removed from the measurement to approximate the audibility of the noise). A preamp with lots of 60Hz hum in the audio signal (caused by AC from the wall getting into the audio signal path) will exhibit a very large increase in S/N ratio (as much as 25dB) when "A"-weighted. If the preamplifier's S/N ratio is referenced to 1V output instead of 0.5V, subtract 6dB from the S/N to find the equivalent 0.5V figure.

Preamps rarely have frequency-response errors; virtually all modern preamps have flat response from below 20Hz to above 50kHz. Phono stages, however, often have errors in their RIAA response, producing the equivalents of frequency-response errors. These errors are caused primarily by capacitor tolerances in the RIAA circuit. An example of a preamp with RIAA errors is shown in [Fig.16-1a](#). The solid line is the left channel, the dotted line the right channel. A phono preamplifier with no RIAA errors would produce a perfectly flat line. The rolloff in the upper midrange and treble seen in [Fig.16-1a](#) will make this phono stage sound dark, closed-in, and lacking immediacy. A phono preamp with more accurate RIAA equalization is shown in [Fig.16-1b](#).





Fig.16-1a & b Examples of poor (left) and typical (right) RIAA accuracy (Courtesy *Stereophile*)

A preamplifier's *total harmonic distortion* (THD) is often measured and expressed as THD+N (total harmonic distortion plus noise). This expression reflects the technique that measures the sum of harmonic distortion and noise. Here's how THD+N is measured: The preamp under test is driven by a swept sine wave signal, a band-reject filter removes the test signal, and whatever is left (distortion and noise) is plotted as a function of frequency. Note that very-low-distortion preamps can have a higher THD+N figure than their intrinsic harmonic distortion; the noise, rather than distortion, dominates the measurement.

The THD+N vs. frequency plot of a typical preamplifier is shown in Fig.16-2. The vertical scale is THD+N as a percent of the signal. The upper pair of traces is measured through the phono inputs, the lower pair through the line inputs. The apparently higher distortion through the phono inputs is undoubtedly caused by its higher noise level, not increased distortion.

Interchannel crosstalk, also called *channel separation*, is a measure of how much the signal in one channel leaks into the other channel. The higher the channel-separation figure, the better, with 100dB being excellent and 50dB being on the poor side. A crosstalk specification should include the measurement frequency; crosstalk usually increases as frequency rises. Crosstalk will generally increase at the rate of 6dB per octave, due to capacitive coupling between channels.

Stereophile Preamp THD+N, line stage (bottom traces) and phono stage




Fig.16-2 THD+N measurement of a typical preamplifier (Courtesy Stereophile)

When expressing this phenomenon as *crosstalk*, the *lower* the specified crosstalk, the *better*; this figure indicates how **much** signal has leaked into the other channel. Conversely, when expressing *channel separation*, the *higher* the figure, the *better*; the number indicates the degree of isolation between channels. Both terms, however, describe the same phenomenon. Fig.16-3a shows mediocre crosstalk performance; Fig.16-3b shows excellent crosstalk performance. The measurement is made by driving one channel of the preamplifier with a swept sinewave at a reference output level (the 0dB reference on the graph's vertical scale), then plotting the undriven channel's output signal as a function of frequency. The measurement is repeated with the channels switched, and both traces are combined on the same graph. In the example of poor crosstalk performance, the crosstalk in the undriven left channel was 60dB below that of the driven right channel at 1kHz. In Fig.16-3b, the signal in the undriven channel was down by 100dB at 1kHz.





Fig.16-3a & b Examples of poor (left) and excellent (right) crosstalk (Courtesy *Stereophile*)

Channel balance is a measure of how closely the output levels of the two channels are matched. Some preamplifiers have slightly more gain in one channel than the other, creating a channel imbalance. Volume controls can also introduce channel imbalance, particularly at the low end (counterclockwise) of their range. Channel balance is often measured at three volume-control positions, and should be within 0.2dB at all volume settings. Large channel imbalances can cause the center image to shift to one side. Tubed preamps sharing a single dual tube for the left and right channels are particularly susceptible to channel imbalances; the two elements within the tube may have slightly different gains. Fortunately, the problem can be corrected simply by replacing the tube with one with matched gain.

Preamplifier specifications should be used to ensure electrical compatibility with other components, not as indicators of sound quality. Many superb-sounding preamps measure poorly, and some preamps that look great on paper disappoint in the listening room. Let your ears—not the specification sheet—be your guide in choosing a preamplifier.

Power Amplifiers

Power amplifiers share many specifications (signal/noise ratio, crosstalk, etc.) with other components, yet have some unique specifications and measurements.



Power amplifiers typically have input impedances of between 20 and 150k ohms, with 47k ohms being a standard for unbalanced inputs. This is a high enough value to ensure that the power amplifier's input stage won't load down the preamplifier's output and put a strain on the preamp's output stage.

A power amplifier's output impedance is usually a very low value—between 0.05 and 0.5 ohm for solid-state amplifiers. Tubed power amplifiers have a higher output impedance—about 0.5 ohm to as high as 2 ohms. Single-ended tubed amplifiers can have output impedances as high as 3 ohms. Think of output impedance as a resistor in series with the output transistors (or output transformer, in a tubed amplifier) connected to the load. Low output impedance is a desirable design goal; when the power amplifier's output impedance becomes a significant fraction of the speaker's input impedance, the speaker and amplifier interact. The result can be peaks and dips in the amplifier's frequency response caused by the speaker's impedance change with frequency. This is a serious consideration when contemplating a single-ended tubed amplifier; the amplifier will have a different tonal balance when driving different speakers. A high output impedance is also characterized by soft bass, lack of bass extension and dynamics, and less bass control. A power amplifier with a high output impedance is less able to control woofer motions. Output impedance is related to an amplifier's damping factor—its ability to control cone motion in the loudspeaker. The higher the damping factor, the better. Note that damping factor is specified at the amplifier's output terminals; resistance in loudspeaker cables can dramatically decrease the effective damping factor, particularly in long cable runs.

Gain refers to the ratio between the input and output voltages. Gain can be specified as the ratio between the voltages, or as a figure in decibels (dB). Typical gain is about 26dB, or a ratio of 20:1 (i.e., an input voltage of 1V will produce 20V at the output). Input sensitivity is the input voltage required to drive the amplifier to clipping. A typical value is between 0.775V and 2V, with about 1.2V being standard. If the power amplifier has a very high input sensitivity (less than 0.5V), the preamplifier's volume control will be extremely insensitive—you'll get plenty of volume with the knob barely turned up. Conversely, a low-output preamp

driving a low-sensitivity power amplifier (2V, say) will require the preamp's volume control to be turned up very high to get adequate volume.

We've talked a lot about maximum output power; now let's see how it's measured. With a dummy load (a huge power resistor) connected to the amplifier's output, the power amplifier is driven with a gradually increasing voltage. Distortion in the output is monitored and plotted as a function of output power. When the amplifier can deliver no more power, the distortion shoots up instantaneously as the waveform "clips." The clipping point, defined as 1% THD, is the amplifier's maximum output power. This test is performed with the amplifier driving loads of 8, 4, and 2 ohms. The results are combined and plotted on a single graph (Fig.16-4a). The farther to the right the "knee" (bend) in the distortion trace, the higher the output power at clipping. Note that some plots of THD vs. output power have only two traces, at 8 and 4 ohms. These amplifiers cannot be tested into 2 ohms at full output power: the power-supply fuses blow, or the amplifier simply shuts down.



Figs.16-4 a and b The power output vs. THD curves show an amplifier's maximum output power, and how well it can increase its power into low-impedance loads. (Courtesy *Stereophile*)

By combining the three traces on one graph, we can quickly assess the amplifier's ability to increase its output power into low impedances. The greater the distance between the curves' knees, the greater the amplifier's ability to drive current into low-impedance loads. Fig.16-4b shows the measurement of an

amplifier with very little ability to deliver current into a 4 ohm load: It produced 15.6dBW into 8 ohms, but only 12.5dBW into 4 ohms.

By contrast, Fig.16-4a shows an amplifier behaving as a virtually perfect voltage source, maintaining its voltage and doubling its current output each time the load impedance is halved. This graph, the result of an actual measurement, reveals exemplary—and extraordinary—performance. The performance would have been even better had the AC line voltage not drooped under the heavy load, decreasing the raw power available to the amplifier—the amplifier pulled so much current out of the wall that the AC line voltage dropped from 118V to 106V. This amplifier put out 325W into 8 ohms (25.1dBW), 635W into 4 ohms (25dBW), 1066W into 2 ohms (24.3dBW), and 1548W into 1 ohm (22.9dBW). Moreover, it was perfectly stable driving this tremendous amount of current into such a low impedance.

These methods are used for most tubed and solid-state push-pull amplifiers. Single-ended tubed amplifiers produce so much distortion that a different criterion must be used to determine the clipping point. Rather than 1% THD as the point at which we say the amplifier has run out of power, the clipping point of a single-ended amplifier, is often said to be 10%. That's because a single-ended tubed amplifier may produce 2 or 3% THD when operated at just a few watts of output power.

Another way to look at a power amplifier's output capability is the power-bandwidth test, which determines the widest bandwidth over which the amplifier can maintain full output power. The amplifier is driven to just below clipping, and the output power is plotted against frequency.

How much peak power above its continuous power rating an amplifier can deliver on transient signals (musical peaks) is called its *dynamic headroom*, expressed in dB. A power amplifier rated at 100Wpc (RMS) that can deliver 200Wpc for brief periods (measured in milliseconds) without severe distortion is said to have 3dB of dynamic headroom. You may remember from earlier in the chapter that doubling the power represents an increase in volume of 3dB. An amplifier's dynamic-headroom specification says a lot about its subjective ability to play loudly.

In addition to plotting THD against output level to find the clipping point, a power amplifier's distortion can be measured at low signal levels and plotted as a function of frequency. Fig.16-5 is such a graph. The lower trace is THD+N (THD plus Noise) measured at 1W into 8 ohms. The next higher trace is at 2W into 4 ohms, and the upper trace (highest distortion) is with 4W driving a 2 ohm load.

The lower the distortion and noise, the better, with less than 0.2% constituting good performance.




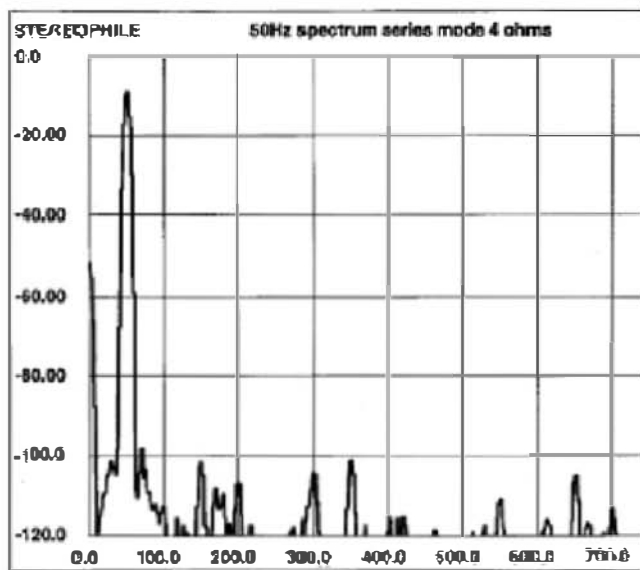

Fig.16-5 THD vs. frequency measurements show how an amplifier's distortion changes into low-impedance loads. (Courtesy *Stereophile*)

Although knowing how much total distortion a power amplifier generates is important, an examination of the harmonic-distortion spectrum is more instructive. We want to know which harmonics are present; some harmonic products are much more audible and annoying than others. As stated earlier, in the discussion of tubed power amplifiers, 10% second-harmonic distortion is less musically objectionable than 0.5% seventh-harmonic distortion. Thus, rather than use a single THD figure, we need to characterize the distortion spectrum.

To do this, the power amplifier is driven at two-thirds rated power with a 50Hz sinewave, and a Fast Fourier Transform (FFT) is performed on the output signal. The FFT captures the amplifier's output signal and transforms it into its constituent parts. (Specifically, the FFT can convert a time-domain signal into the frequency domain.) An FFT produces a plot showing energy vs. frequency, as shown in Fig.16-6a. We can see that this particular amplifier generates an entire

series of harmonic-distortion products. The distortion spectra in Fig.16-6b, however, show a much lower level of distortion products.

Although harmonic distortion has been touted as the definitive performance parameter of power amplifiers and receivers, this specification in itself isn't that significant a criterion in the evaluation of power amplifiers. Many amplifiers with 3% distortion sound just fine, while some with only 0.01% sound awful. As described in Chapter 1, THD can be reduced by increasing the amount of negative feedback in the amplifier. Unfortunately, negative feedback introduces other problems that make the amplifier sound worse. Although it's interesting to know how much harmonic distortion an amplifier produces and the spectral distribution of those distortion products, I wouldn't let THD figures influence my purchasing decision.



Figs.16-6a and b An FFT of an amplifier's output shows the amplifier's individual harmonic-distortion components. (Courtesy Stereophile)

Intermodulation distortion (IMD) is distortion that occurs when two signals of different frequencies are amplified simultaneously—something that happens all the time in music. IMD creates sum-and-difference signals related to the frequencies being amplified. For example, if we drive a power amplifier with 10kHz and 1kHz sinewaves, the amplifier produces an output signal containing the original test tones at 10kHz and 1kHz, but also intermodulation products at 11kHz (the sum of the two signals) and at 9kHz (the difference between the two signals). Though the 9kHz and 11kHz sum-and-difference IMD products are

usually the highest in amplitude, IMD creates an infinite series of distortion products. Intermodulation distortion is expressed as a percentage of the IMD products in relation to the test-signal amplitude, with typical values falling between 0.01% and 0.2%. IMD test signals are often 60Hz and 7kHz combined at equal amplitude, or with the higher-frequency tone at one-quarter the amplitude (-12.04dB), or 19kHz and 20kHz at equal amplitude. To test with the latter signals, the amplifier is driven at two-thirds rated power and an FFT is performed on the amplifier's output. We look for the 1kHz difference signal (20kHz minus 19kHz); the lower its amplitude, the better. This 1kHz intermodulation product is usually between 55dB down (i.e.; 55dB below the test-signal amplitude) and 75dB down. We also sometimes see a family of sideband signals around the test tones.

Some of these measurements are unique to power amplifiers. Other, more universal measurements include those of frequency response, crosstalk, and signal/noise ratio. Except for single-ended amplifiers, however, it's rare for a modern amplifier to depart from flat frequency response. Note that the sonic characteristics of "bright," "extended," "forward," etc. are *not* functions of the amplifier's frequency response.

Finally, a power amplifier is said to be *inverting* if it reverses absolute polarity between its input and output. A non-inverting amplifier maintains the same polarity between input and output. In an inverting amplifier, the positive- and negative-going parts of the wave are flipped over: positive at the input becomes negative at the output. This can slightly improve the sound of some recordings and degrade others. Because most gain stages invert absolute polarity, an inverting amplifier usually has an odd number of gain stages, while a non-inverting amplifier has an even number; one isn't necessarily better than the other. (See "Absolute Polarity" in [Appendix A](#).)

Loudspeakers

There is more correlation between a loudspeaker's measured performance and its sound quality than there is with any other component in an audio system. A loudspeaker that measures as having a treble rolloff will likely sound dull; a speaker with a measured rising top end will likely sound bright and forward. Having said that, however, I must stress that you can't judge a speaker's performance by reading its specifications or measurements. Although good measured performance is valued, the speaker's musical abilities in the listening room are what count. To find that out, we must listen.

Following is an explanation of each of the measurements found in some speaker reviews or manufacturers' specification sheets.

Impedance is the electrical resistance the speaker presents to the power amp. As described in detail in [Appendix B](#), a loudspeaker's impedance is a combination of its simple DC resistance, inductive reactance, and capacitive reactance. The higher the impedance, the easier it is for the amplifier to drive. A low-impedance loudspeaker requires the amplifier to drive more current through the loudspeaker, which puts more stress on the amplifier.

Even though a loudspeaker's impedance may be specified as "6 ohms," the actual impedance almost always varies with frequency. [Fig.16-7a](#) shows an impedance curve (solid line) for a loudspeaker of rather low impedance. The impedance dips to less than 4 ohms (the "40.0m" horizontal division) through the bass and most of the midrange. The impedance minima is just over 1 ohm at 30Hz—an extremely low value.



The dotted trace in the impedance plot shows the speaker's phase angle, which indicates how inductive or capacitive the speaker "looks" to the amplifier. If the loudspeaker presented a pure resistance (no capacitance or inductance), the phase-angle plot would be a straight line. The farther away from the center line, the more reactive the speaker is at that frequency. (Phase angle is discussed in detail in [Appendix B](#).)

A combination of high reactance and low impedance is especially difficult for a power amp to drive. The severe phase-angle swings seen in [Fig.16-7a](#), coupled with the speaker's low impedance, will be more demanding of the power amplifier. This doesn't mean that the loudspeaker won't work, only that it will work better with a power amp that can deliver a substantial amount of current into low-impedance loads. For example, a power amp that nearly doubles its

power into a 4 ohm load compared to an 8 ohm load would be a better choice than one that barely increases its power into 4 ohms.

A much easier load to drive is illustrated in Fig.16-7b. The impedance is much higher over a wider band, and the phase-angle swings are less severe. The impedance never dips below 6 ohms, and is well above that figure over wide bands. This loudspeaker will demand much less current from the power amplifier.

A loudspeaker's impedance magnitude and phase angle must be considered in relation to its sensitivity. A speaker with a high sensitivity is much less demanding of a power amp than one with a low sensitivity, regardless of the loudspeaker's impedance. The combination of low sensitivity, low impedance, and severe phase angle requires an amplifier of the highest quality to achieve musically acceptable results.



Figs.16-7a and b A loudspeaker's impedance magnitude (solid line) and phase angle (dotted line) indicate how easy or difficult it is for a power amplifier to drive. (Courtesy *Stereophile*)

Sensitivity can be expressed as a sound-pressure level with 1 watt input measured from 1 meter away. A loudspeaker's sensitivity may be expressed as 88dB/1W/1m. High sensitivity is generally considered any value above 90dB/1W/1m, while low sensitivity is a value below 85dB/1W/1m.

Note that sensitivity can also be expressed with a drive signal of 2.83V, which corresponds to 1W into an 8 ohm load. Loudspeaker manufacturers can cheat on

sensitivity figures by using the 2.83V figure into a 4 ohm speaker, which increases the speaker's sensitivity rating by 3dB. The speaker isn't 3dB more sensitive, it's just drawing twice as much current from the power amp at the same 2.83V. If you see the 2.83V figure, make sure the impedance is 8 ohms. Note that the term *efficiency* is often used incorrectly in place of *sensitivity*. Technically, efficiency is the percentage of electrical power converted by the loudspeaker into acoustical power.

Measuring a speaker's impedance and phase angle can be done electrically on a test bench. Measuring a speaker's acoustical characteristics, however, requires putting a microphone in front of the speaker and driving the loudspeaker with test signals. Consequently, the measured response will be not the speaker's intrinsic characteristics, but the loudspeaker's response as modified by the room in which the measurements are taken. As we know, a speaker's performance is greatly dependent on the room in which it is used. We could end up measuring more of the room's characteristics than the speaker's.

One way of solving this problem is to measure the speaker in an anechoic chamber—a room whose surfaces are covered with highly absorbent material. This material absorbs virtually all the sound energy striking it, reflecting no energy back into the room. (In fact, *anechoic* literally means “no echo.”) Consequently, the measurement microphone “sees” only the speaker's response.

Anechoic chambers are expensive to build and consume lots of real estate. Engineers have therefore developed methods of measuring loudspeakers in a normal room to exclude the room's influence from the measured results. One of these techniques, invented by Richard Heyser, is called time-delay spectrometry (TDS). In TDS measurements, the loudspeaker under test is driven by a swept sinewave. The signal picked up by the measurement microphone is filtered by a narrow filter that allows only a narrow range of frequencies to pass. The frequency the filter allows to pass tracks the swept sinewave driving the loudspeaker, but is delayed in time by a few milliseconds (ms). The filter's delay is exactly the amount of time it takes for sound from the speaker to reach the measurement microphone. The sound reaching the mike is the speaker's direct acoustic output, followed a few milliseconds later by room reflections; or, the same sound after it has been reflected off of the room's walls, floor, and ceiling.

This is where the swept bandpass filter comes in. The filter tracks the sinewave stimulus so that when a reflection reaches the microphone, the filter has already moved past the reflection's frequency and rejects it. The result is the loudspeaker's intrinsic response, unmodified by the room.

A second, more popular technique is implemented in a personal-computer-based hardware and software package called the Maximum Length Sequence System Analyzer (MLSSA). The speaker under test is driven by a stimulus containing a specific pattern of noise sequences. A measurement microphone placed a few feet in front of the speaker picks up this signal, which is compared to the noise-sequence pattern stored in memory. From how the speaker has altered this sequence, the system computes the speaker's impulse response. The period between the impulse and the first reflection is free from reflections and is indistinguishable from an anechoic measurement. The impulse response is then "windowed" to remove the information from the first reflection onward, leaving only the anechoic window.

The impulse response, which is a time-domain signal, is transformed into the frequency domain using Fourier analysis. Fourier analysis is a mathematical technique of converting an impulse response (a time-domain signal) into a frequency response (a frequency-domain signal), or vice versa. The term Fast Fourier Transform (FFT) describes this conversion from one domain to another.

From the loudspeaker's impulse response we can derive the system's frequency response. The result is the speaker's response minus the room's contribution. Note, however, that the anechoic window's length determines the low-frequency resolution of the measurement. The longer the anechoic portion of the impulse response that is transformed into the frequency domain, the lower the frequency that can be accurately measured. When we measure a loudspeaker using this technique, the speaker is placed on a tall stand so that it is halfway between the floor and ceiling, maximizing the anechoic window. The room should therefore be as large as possible.



Despite these techniques, the low-frequency resolution of MLSSA measurements isn't satisfactory unless you have a huge room (a long anechoic window). To better measure a loudspeaker's low-frequency performance, a measurement microphone is placed so that it nearly touches the woofer cone. At this position, the room's contribution is negligible. This separately measured nearfield woofer response is appended to the MLSSA-derived response to provide an accurate wideband measurement.

With that background, let's go through a set of measurements that you might find in a typical loudspeaker review and examine what they mean.

A speaker's impulse response reveals its time behavior, including how time-coherent the loudspeaker is, how much the drivers exhibit ringing, and whether the drivers are connected in the same polarity. I described *time coherence* in the

section on crossovers; a time-coherent speaker has no phase lag between its drivers. That is, they all move in step in response to the input signal. In a speaker that isn't time-coherent, the individual drive units don't all move at the same time in response to the input signal; some take longer to respond.


Fig.16-8a is the impulse response of a loudspeaker that has poor time coherence. Notice the negative-going spike (the tweeter's output), which indicates that the tweeter's output occurs before the other drivers' outputs, and that it is connected in the opposite polarity. The impulse response in Fig.16-8b is much more coherent; all the drivers respond to the impulse in an identical manner. Ideally, the impulse response should have a fast rise-time (indicated by the steepness of the impulse) and minimal driver ringing after the impulse has stopped. Ringing (driver motion after the input signal has stopped) can be seen as squiggles after the impulse.



Figs.16-8a and b A poor impulse response (left) and an excellent impulse response (right). (Courtesy *Stereophile*)

Another way of looking at a speaker's time behavior is the step response, or how the speaker behaves when presented with direct current (DC). Derived from the impulse response, the step response plots the loudspeaker's output vs. time. Fig.16-9a is the step response of the same speaker whose impulse response is shown in Fig.16-8a. The lag between the tweeter (the negative-going spike) and the other drivers is apparent, as well as a slow risetime (the time it takes for the

loudspeaker to reach full amplitude in response to the input signal). For contrast, Fig.16-9b is a virtually perfect step response. The drivers are time-coherent, and the risetime is extremely fast (the step is a straight vertical line). This step response was derived from the impulse response in Fig.16-8b.



Figs.16-9a and b A poor step response (left) and an excellent step response (right). (Courtesy *Stereophile*)

Fig.16-10 is a loudspeaker's overall frequency response as measured over a 30° horizontal window. This simply means that the frequency response was measured at various angles from the speaker's front baffle and averaged to produce the curve. Averaging the response gives a better representation of the listener's subjective response to the loudspeaker; some of that side energy will reach the listener after being reflected from the listening room's sidewalls. This particular speaker has some excess treble energy between 5kHz and 15kHz, which may make it sound tizzy in the treble. Notice how the response drops rapidly above 15kHz. The broad dip between 500Hz and 5kHz may give the loudspeaker a distant rather than a present perspective. The rising bass indicates that the low-frequency alignment may be underdamped, giving the bass a generous but fat quality. In short, these deviations from a flat line indicate colorations in this speaker.




Fig.16-10 Frequency response averaged over a 30° horizontal window (Courtesy *Stereophile*)

Fig.16-11 is a speaker’s “horizontal family” of response curves. These show how a speaker’s response changes as the measurement microphone is moved around the speaker. The frequency-response measurement just described is an average of these individual curves.




Fig.16-11 A horizontal response family. (Courtesy *Stereophile*)

The trace representing the frequency response directly in front of the speaker (on axis) has been normalized to be a straight line so that the other curves show only the difference from the on-axis response. The horizontal family shows the loudspeaker's dispersion, or how it radiates acoustical energy into the room as a function of frequency. This particular loudspeaker shows a typical treble rolloff as the microphone is moved to the sides of the cabinet, but also excess treble energy on axis. These curves suggest that this speaker should be listened to without toe-in; the excess treble on-axis will be reduced if the listener sits slightly off-axis (no toe-in). A speaker that exhibits radical changes in off axis response can sound colored, even though its on-axis response may be fairly flat.

The next measurement (Fig.16-12) shows how the speaker's response changes with listening height. Again, the flat line has been normalized so that only the differences from vertical on-axis response are shown. The on-axis point is usually the tweeter axis, or a typical listening axis (the text accompanying the measurement will usually indicate which). As you can see, the frequency response changes radically with listening height. The speaker has severe peaks and dips in the response (caused by interference effects) above and below the optimal

listening axis. This means the speaker will sound very different depending on listening height. This example is extreme; most speakers don't exhibit such radical response changes along the vertical axis. All speakers, however, evince some change in tonal balance at different listening axes. This measured response underscores the importance of listening height when auditioning a loudspeaker.





Fig.16-12 A vertical response family (Courtesy *Stereophile*)

The ideal speaker would stop generating acoustical energy as soon as the input signal stopped. In practice, some amount of acoustical energy, however small, is always stored in the drivers and enclosure. This energy is radiated over time, like a bell ringing. To look at a speaker's response over time, its frequency response is plotted in discrete slices of time. As described previously, the frequency response is derived from the impulse response using the Fast Fourier Transform. By moving the window in the impulse response on which the FFT is performed, the speaker's frequency response can be derived for different points in time. Specifically, the FFT-derived frequency response is calculated perhaps 100 times over 4 milliseconds (ms) to create the plots of [Figs.16-13a and b](#). This graph is properly called a cumulative spectral-decay plot, but is more commonly called a waterfall plot because of its similarity to cascades of falling water. It shows how the speaker's output changes over time and with frequency. The topmost curve is the speaker's frequency response at zero time—the instant it receives the signal.

The next lower line is the frequency response 40 microseconds (40 μ s) later. The next lower line is 40 μ s after that, and so forth.




Figs.16-13a and b Poor (left) and excellent (right) cumulative spectral decays. (Courtesy *Stereophile*)

By plotting the response over time, we can easily see cabinet resonances, driver ringing, driver breakup, and other problems. A driver resonance is seen in the figure above as a vertical ridge in the plot; the loudspeaker continues to radiate energy at the ridge's frequency long after the impulse has stopped. The perfect loudspeaker would produce a flat line at the top, and the sound it produced would instantly decay to no energy output.

Examples of poor and excellent performance on the cumulative spectral-decay test are illustrated in Figs. 16-13a and b. In Fig.16-13a we can see a ridge at 3729Hz, indicating that the loudspeaker continues to produce acoustic energy at this frequency for 2ms after the drive signal has ceased—a long time in loudspeaker terms. By contrast, Fig.16-13b shows exemplary behavior: the acoustic output drops very quickly (the white space between the top few traces and the lower traces), and almost no ridges are apparent. The one minor resonance, at 2963Hz, is minuscule compared to the resonances seen in Fig.16-13a, which will be audible as a change in timbre when this frequency is excited by part of the music.

The same technique can be used to examine the behavior of loudspeaker cabinets. An accelerometer is attached to the speaker's side panel, which is driven by the MLS signal. The accelerometer's output is drawn in a waterfall plot, which shows the frequency and amplitude of cabinet vibrations. For example, Fig.16-14 shows a speaker enclosure with a severe resonance at 187Hz, seen as the narrow

ridge at that frequency. This loudspeaker stores, then slowly releases, energy at 187Hz. This will undoubtedly color the bass when the speaker is excited by music with energy at this frequency.



-5.38 dB, 184 Hz (47), 0.000 msec (0)

Fig.16-14 Cumulative spectral-decay of the output of an accelerometer attached to a loudspeaker cabinet.. (Courtesy *Stereophile*)

Readers wanting to delve deeper into loudspeaker design are directed to Martin Colloms's definitive reference book on the subject, *High Performance Loudspeakers* (Halsted Press, ISBN 0-470-21721-9).

Digital Components

There are many tests and measurements that quantify a DAC's technical performance, but few—if any—tell us how the DAC sounds. Nonetheless, good measured performance indicates that the DAC is well engineered, and is more likely to have good sonic performance as well.

Let's start with a DAC's maximum output level. (I'll use the term DAC, although I'm also referring to CD players—the measurements are the same.) The standard analog output voltage is 2VRMS, with individual units varying between 1.75V and a whopping 7.2V. Most CD players and DACs put out between 2.2V and 3.5V. Note that this value indicates the highest RMS output voltage possible from the player, measured by playing a full-scale, 1kHz sine wave.

Knowledge of a DAC's output level is useful in several ways. First, a DAC of very high output wouldn't be ideal for driving a preamplifier that had lots of gain. Although it's unlikely that the DAC would overload the preamp's input, you would end up using the preamp's volume control only at the very low end of its range, where most volume controls exhibit their greatest channel imbalances (i.e., one channel is louder than the other). In addition, setting a "just right" volume is more difficult with high-output DACs and high-gain line-stage preamplifiers; the volume control becomes overly touchy to very small movements. This is particularly true of preamps with detent level steps instead of a continuously variable volume control. A high output level is an asset when driving a passive level control, but a drawback with a high-gain line-stage preamplifier.

Another important measurement to consider when determining if a particular DAC or CD player is a good candidate for a passive level control is output impedance. If a passive level control is used between a DAC and power amplifier, the burden of driving the cables and power amplifier falls on the DAC's output stage. If the DAC has a high output impedance, dynamics can become compressed and the bass may get mushy. In extreme cases, a DAC with a very high output impedance can even become current-limited and flatten the waveform peaks if asked to drive a low-input-impedance preamplifier. DACs with the highest output impedances have an impedance of nearly 2k ohms at 20kHz and 1.5k ohms at 1kHz. This high output impedance causes the DAC's output stage to clip when asked to drive impedances of less than about 15k ohms.

Further, the output impedance of the passive level control itself must be added to the equation. The passive control's output impedance (which varies

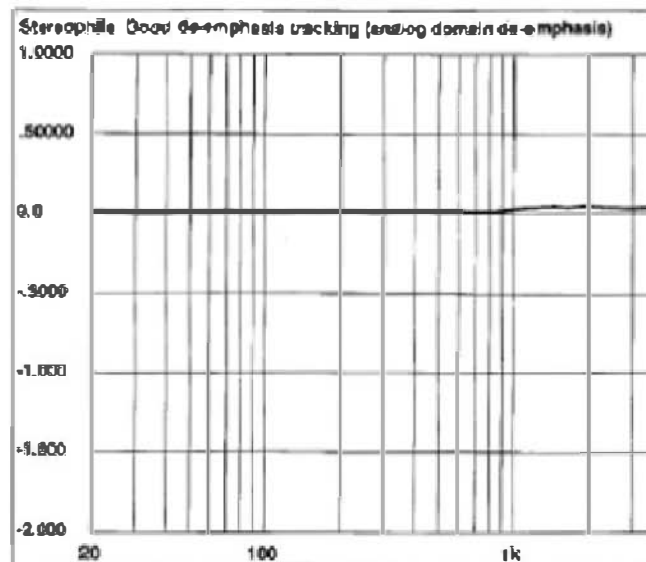
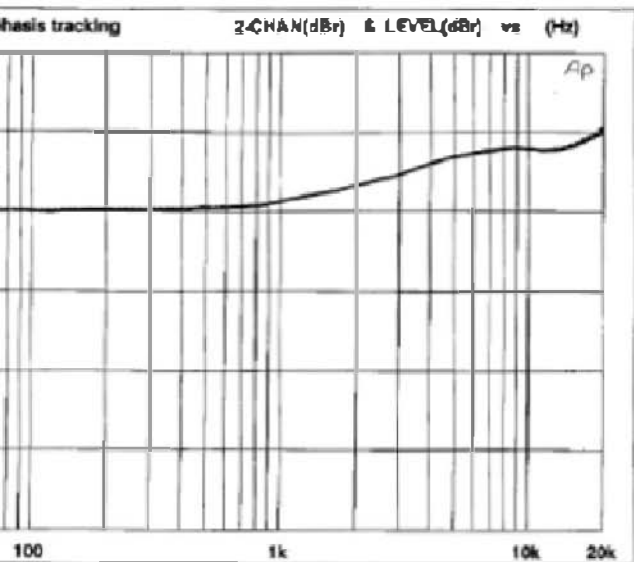
depending on which resistor is switched in, or the position of the potentiometer) must be added to the driving component's source impedance to find the total source impedance driving the power amplifier.

DAC output impedances vary from less than 1 ohm to 5.6k ohms (the highest value I've measured in a product). Generally, if the output impedance is less than a few hundred ohms, the processor should have no trouble driving a passive level control. Typical values are between 50 and 500 ohms.

Bottom line: Don't drive a low-input-impedance power amplifier (less than about 20k ohms) with a preamp that functions as passive level control, with a high-output-impedance (greater than about 800 ohms) DAC through long, high-capacitance interconnects. These aren't hard-and-fast numbers but general guidelines. [Appendix B](#) includes a full analysis of what can happen with a passive level control and high-output-impedance sources.

The next measurement, frequency response, doesn't reveal a processor's overall tonal balance, as it often does with speakers, but is included in specifications and measurements for the sake of completeness. Most processors are down a few tenths of a dB at 20kHz, the result of rolloff in the digital filter.

More revealing, however, is a CD player's or DAC's de-emphasis error. As previously described, some CDs have had their trebles boosted (emphasized) during recording, and carry a flag that tells the CD player to switch in its de-emphasis circuit to restore flat response. These de-emphasis circuits are often inaccurate (due to resistor and capacitor tolerances), producing frequency-response irregularities when the player is decoding pre-emphasized discs. De-emphasis error should be less than $\pm 0.1\text{dB}$ across the audioband. De-emphasis errors can cause tonal-balance irregularities (tizzy upper treble, for example) when playing pre-emphasized CDs. De-emphasis errors of 0.2dB over an octave of bandwidth are audible. Poor and good de-emphasis accuracy is shown in [Figs. 16-15a and b](#). Both of these measurements were made on digital processors using analog de-emphasis circuits; processors with digital-domain de-emphasis have no errors. The positive error in [Fig.16-15a](#) will make the processor sound bright when playing pre-emphasized CDs.





Figs.16-15a and b De-emphasis errors (left) and accurate de-emphasis tracking (right) (Courtesy *Stereophile*)



Interchannel crosstalk, also called *channel separation*, is a measurement of how well the left and right channels are isolated from one another. Ideally, when a DAC is fed a right-channel signal, that signal should not appear in the left channel, and vice versa. The lower the crosstalk (or, put another way, the greater the channel separation), the better. Many DACs (and other products) have decreasing channel separation as frequency increases, a result of capacitive coupling between channels. Typical CD players and DACs have about 90dB of channel separation at 1kHz, this decreasing to about 70dB at 20kHz. Although there is no correlation between low crosstalk and soundstage width (an intuitive link), high channel separation indicates good overall engineering. Mediocre and good crosstalk performance are shown in Figs. 16-16a and b. At 1kHz, the processor shown in Fig.16-16a has 67dB of channel separation, compared with the 112dB of channel separation at 1kHz measured in the processor shown in Fig.16-16b.

A very revealing test of a processor's ability to re-create low-level signals looks at the processor's output vs. frequency (spectral analysis) when the processor under test is driven with the digital code representing a -90dB, dithered 1kHz sinewave. A spectral analysis plots energy level vs. frequency. Peaks in the trace at 60Hz, 120Hz, 180Hz, and related frequencies indicate that power-supply noise is getting into the audio circuitry. A spectral analysis can also hint at how well the unit's DACs are performing; the signal should peak at the -90dB horizontal

division, and the left and right traces should overlap. This indicates that the DACs are linear and perform similarly between channels. The spectral analysis also shows the processor's noise level, and reveals if there is any spurious energy in the audioband. Figs. 16-17a and b respectively show poor and good performance on this test. In the poor-performing DAC, we can see power-supply noise in the audio signal at 60Hz and harmonics of 60Hz, linearity error (the traces don't peak at the -90dB horizontal division), no overlap of the traces (indicating that the DACs are not well matched between channels), and a higher overall noise level. The performance shown in Fig.16-17b is exemplary.



Figs. 16-16a and b Poor (left) and good (right) crosstalk performance (Courtesy *Stereophile*)



Figs. 16-17a and b Poor (left) and good (right) performance on the spectral-analysis test (Courtesy *Stereophile*)

The same type of spectral analysis, but with the DAC driven by “digital silence” (all zeros), is also performed and plotted to 200kHz. This test can reveal the presence of spurious junk in or above the audioband, and DAC artifacts such as idle tones (seen as a peak in the trace). Ideally, the trace should be a smooth line sloping gently upward, with no peaks at the power-line frequencies and no spikes in the audioband. The effects of noise shaping in 1-bit converters can also be seen in this measurement. As described earlier in this chapter, noise shaping is a technique that pushes noise from the audioband into the ultrasonic band. [Fig.16-18](#) shows the effects of noise shaping as revealed by this wideband, 1/3-octave spectral analysis: the noise floor rises sharply above 10kHz. We can also see some peaks in the left-channel trace (the solid line) at 18kHz and 25kHz that may be DAC artifacts.

Low-level linearity was touted a few years back as the definitive measurement in digital processors. While good linearity is a high engineering goal, it doesn’t guarantee good sound. Understanding linearity is easy: it’s how well the processor’s analog output level tracks the digital input level. If we input the code representing a -88dBFS signal, for example, the analog output of the test tone should be exactly -88dB (88dB below full scale). Any difference between digital input level and analog output level is linearity error. Linearity error is expressed as the difference in dB between the digital input level and the analog output level

at a specific digital level (with any gain error normalized out of the measurement). For example, we can say that a particular converter has a 3dB positive error at -90dBFS . This means that when the DAC is driven by the digital code representing a -90dBFS signal, the DAC's analog output is actually -87dB , or 3dB higher than the digital input level. Note that the digital input level refers to the signal levels represented by the value of the ones and zeros in the binary code, not the input voltage to the DAC's digital input, which remains constant.






Fig.16-18 A wideband spectral analysis reveals the effects of noise shaping in 1-bit converters. (Courtesy *Stereophile*)

A comparison of mediocre and excellent linearity is shown in Fig.16-19. A perfect processor would produce a ruler-flat line, with no deviation above or below the zero horizontal division. The linearity measurement in Fig.16-19b comes very close to this ideal. Conversely, the measurement in Fig.16-19a has an error of -2.6dB in the right channel (dotted trace) and an error of -4dB in the left channel (solid trace) at a signal level of -90dB .



Figs. 16-19a and b Mediocre (left) and excellent (right) linearity (Courtesy *Stereophile*)



If the linearity trace dips below the center line, this indicates a negative linearity error. If it rises above the center line, the processor has a positive linearity error (the analog output is higher than it should be). The digital input level is read across the bottom, with the highest level at the graph's right-hand side. Some processors with excellent low-level linearity appear to have a positive error at very low signal levels (-115dB). This is actually noise swamping the DAC's output. The practical lower limit of measuring linearity is about -112dBFS (in quiet processors). Linearity is determined by the digital-to-analog converter chip inside the DAC.

Another test is the FFT-derived spectral analysis of the DAC's output when decoding a full-scale mix of 19kHz and 20kHz tones. This test reveals intermodulation components such as the 1kHz difference tone (20kHz minus 19kHz), and interactions between the sampling frequency (44.1kHz) and the test signals.

As with the spectral analysis we saw earlier of the -90dB , 1kHz sinewave, this spectral analysis also plots energy against frequency. But this technique does it a little differently. In an FFT-derived spectral analysis, the DAC's output is captured in the time domain, creating a waveform that is then transformed into the frequency domain using the FFT.

The result of this technique is shown in Figs. 16-20a and b, the FFT-derived spectral analyses of two processors' outputs when decoding this unusual test

signal. Ideally, there should be no spikes in the noise floor; Fig.16-20b shows exemplary performance on this test. Note the greater number of spikes in Fig.16-20a. The higher the spikes' amplitude and the greater their number, the worse the intermodulation distortion performance. The trace in Fig.16-20b reveals a 1kHz difference component, but very few other intermodulation products.





Figs.16-20a and b Some processors produce high intermodulation distortion (left); others generate very little IMD (right). (Courtesy *Stereophile*)

DACs run into problems when converting very-low-level signals to analog. One way to examine a DAC's behavior with small signals is to capture its analog output waveform when reproducing a -90dB , undithered 1kHz sinewave. This test signal produces three quantization steps: 0, +1, and -1 . The three levels should be of equal amplitude, and the signal should be symmetrical around the center horizontal division. A poor-quality reproduction of this signal is shown in Fig.16-21a, a high-quality one in Fig.16-21b. Note the absence of noise in Fig.16-21b, and the three clearly defined quantization steps. By contrast, Fig.16-21a is so noisy it's difficult to make out the sinewave's shape.

A DAC's reproduction of a 1kHz, full-scale squarewave (not shown) reveals its time-domain behavior. Most measurements show the overshoot and ringing typical of the digital filters used in most processors and CD players. The digital filter determines the waveform shape; the time distortion occurs in the digital

filter. Some digital processors with custom DSP-based filters don't suffer from this type of distortion.



Figs.16-21a and b Comparison of a noisy and poorly shaped 1kHz -90dB undithered sine wave (left), and nearly perfect reproduction (right). (Courtesy *Stereophile*)

The amount of DC at the DAC's analog outputs is usually measured and noted. The lower the DC level, the better. DC may cause a thump through the system when the preamplifier's input selector is switched to select the DAC's output. If the preamplifier and power amplifier are both direct-coupled (no DC blocking capacitors in the signal path) and use no DC servo correction, the DC could be amplified and appear at the loudspeaker terminals—not a good condition. Low DC—less than, say, 20mV—is therefore desirable.



The measurements section of a digital product review will often mention that the DAC is polarity-inverting or non-inverting. A non-inverting DAC's analog output will be of the same absolute polarity as the digital data on the CD. Most processors don't invert polarity, but if one does, that's no reason for concern. Without knowing the absolute polarity of the recording and the rest of your system, it's a fifty-fifty chance that the absolute polarity will be wrong. Knowing that a DAC inverts absolute polarity is useful when comparing DACs. For the listening comparisons to be valid, all DACs under audition should maintain the same polarity.

Measuring a DAC's word-clock jitter (inaccuracy in the crucial timing signal) was first made possible in 1992 by an instrument called the LIM Detector,

developed by Ed Meitner, of Museatex. LIM stands for Logic Induced Modulation. Although other designers have looked at word-clock jitter, the LIM Detector was the first commercially available jitter-analysis tool. When used with an FFT machine and an RMS-reading voltmeter, the LIM Detector provides the digital processor's overall RMS jitter level over a specified bandwidth and a spectral analysis of that jitter.

To use the LIM Detector, the word-clock pin on the DAC is found, and a probe is connected between the word clock and the LIM Detector. The LIM Detector's output is only the jitter component of that word clock. A wide range of word-clock frequencies is accommodated. The LIM Detector output is fed to the input of an FFT analyzer, which displays the jitter's spectral content. The analyzer's RMS meter displays an overall RMS voltage, which is converted to jitter in picoseconds (ps). We thus end up with a number indicating the overall amount of jitter on the digital processor's word clock, and FFTs showing the jitter's spectral distribution.

The FFT in [Fig.16-22a](#) shows a jitter spectrum in which the jitter is periodic; that is, the jitter energy is concentrated at specific frequencies, seen as spikes in the plot. Where you see a spike, it means the digital processor's word clock is exhibits jitter at that frequency. Note that the spikes occur at 1kHz and multiples of 1kHz; the test tone driving the digital processor is a 1kHz sinewave. As described earlier in this chapter, jitter is often correlated with the audio signal—something we can readily see in [Fig.16-22a](#). The RMS jitter level, measured over a 400Hz–20kHz bandwidth, is 350ps.



Figs.16-22a and b High word-clock jitter in a digital processor (left), and low jitter with few periodic jitter components (right). (Courtesy Stereophile)

For contrast, Fig.16-22b is the jitter spectrum of a very-low-jitter DAC. Not only is the overall RMS level lower (less than 80ps), but the jitter spectrum is nearly free from signal-correlated jitter. In short, the lower the trace on the FFT, and the smoother the line (lack of spikes), the better the DAC's jitter performance.

This method of measuring jitter is direct in the sense that we're looking at clock jitter *where* it matters, but indirect because it doesn't tell us the effect of that jitter on the analog output (*what* really matters). More recently, measurement systems have been developed that measure the amount of clock jitter by analyzing the DAC's analog output. These analyzers better correlate with listening impressions. Moreover, other factors influence how much jitter is actually present at the conversion to analog; many DAC chips add jitter to the clock inside the chip where it is impossible to measure with the LIM technique.

Few measurements are available that even attempt to characterize a transport's performance. One of them, however, can be made by anyone, and gauges how well the transport's tracking and error-correction systems work. All you need is the *Pierre Verany Digital Test CD* (Pierre Verany PV.788031/788032). Disc 2 of this two-CD set contains a section in which the spiral track of pits is interrupted by blank areas (no pits). This simulates data dropouts caused by poor-quality CDs, dirt, scratches, or other media contamination. A continuous

tone is recorded on the tracks. The blank areas gradually increase in length, and each increase represented by a different track number. The higher the track number in this series, the longer the dropout. By listening for when the test tone is interrupted and noting the track number, the dropout length the player can track and correct can be determined. The higher the track number the transport will play flawlessly, the better. Most transports will play up to track 33. Skip-free reproduction of track 37 or higher is excellent.

I must reiterate that a DAC's bench performance doesn't predict its sound quality. There is no way to predict a component's musical character from looking at its test results.

Appendix A: Sound and Hearing

Introduction: What is Sound?

Sound is a series of physical disturbances in a medium such as air. When an object vibrates, it sends out a series of waves that propagate through the air. These waves are composed of fluctuations in air pressure above and below the normal atmospheric pressure of 14.7 pounds per square inch. When a loudspeaker cone moves forward, for example, it compresses the air in front of it, creating an area in which the air molecules are denser than normal atmospheric pressure. This portion of the sound wave is called a *compression*.

When the loudspeaker cone moves backward, it creates an area in front of it in which the air molecules are less dense than atmospheric pressure. This portion of the sound wave is called a *rarefaction*. Sound is made up of a series of alternating compressions and rarefactions moving through the air.

Sound waves are transmitted through the air by passing on the moving molecules' momentum to adjacent molecules. The original compressed molecules return to their original positions because of the elastic properties of air. The air molecules don't move very far; instead, their "bumping into each other" is what transmits sound energy.

When these compressions and rarefactions strike our eardrums, we perceive the phenomenon as sound. The greater the change in air pressure above and below normal atmospheric pressure, the greater the *amplitude* of the sound.

Because most objects vibrate with a periodic back-and-forth motion, or oscillation, most sound waves (and nearly all musical sounds) have a periodic repetition. The sound wave is a replica of the object's motion. Consequently, sound waves have a regular, periodic pattern of compressions and rarefactions.

This periodic pattern is illustrated in [Fig.A-1](#). The loudspeaker at one end of the tube of air is driven by a periodic signal, such as that made by the cyclic back-and-forth motion of a tuning fork. [Fig.A-1](#) (b) shows the movement of the loudspeaker diaphragm as a function of time. When the cone moves forward, a compression is created; when it moves backward, a rarefaction is formed. This

pattern moves down the tube at 1130 feet per second, the speed of sound in air. The resulting pattern of pressure change, shown in Fig.A-1 (c), is called a *sinewave*, the most basic periodic repetition. Fig.A-1 (c) shows the instantaneous pressure distribution over the tube's length. Fig.A-1 (d) represents the change in atmospheric pressure as a function of time at some point along the tube.

Another way of relating periodic motion to a sinewave is to imagine a tuning fork with a pen attached to it, with the pen touching a moving strip of paper. As the paper moves under the pen, the tuning fork's oscillations trace a sinewave on the paper.






Fig.A-1 The relationship between a loudspeaker moving back and forth, pressure distribution in a tube, wavelength, and period (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988)

Period and Frequency

The time it takes to complete one cycle of a sinewave (indicated by the “T” in A-1 (b)) is called its *period*. The *frequency* of a sinewave is the number of periods per second. A sound’s frequency largely determines the sound’s *pitch*, our subjective impression of frequency.

The unit for expressing frequency is the *Hertz*, abbreviated *Hz*. This unit of measurement is named after Heinrich Hertz, the German physicist who established that light and heat are electromagnetic radiations. The use of the unit Hz is relatively recent; the older unit of frequency was the more descriptive *cycles per second*, or *cps*.

Frequency is the reciprocal of time, expressed in the equation $F = 1/T$ where F is the frequency in Hz and T is the wave’s period in seconds. For example, a sinewave with a period of 0.01 seconds (one one-hundredth of a second) has a frequency of 100 cycles per second, or 100 Hertz (100Hz). Because each period occupies one one-hundredth of a second, there are one hundred periods in a second, or 100Hz.

A sinewave with a period of 0.001 seconds (one one-thousandth of a second, or one millisecond) has a frequency of 1000Hz. The shorter the period, the higher the frequency.

Wavelength

If a wave's period is the *time* it takes to complete one cycle, a wave's *wavelength* is the *distance* between successive cycles. Wavelength, represented by the Greek letter lambda (λ) is shown in [Fig.A-1 \(c\)](#).

A sound's wavelength in air is directly related to the frequency: the shorter the wavelength, the higher the frequency. To calculate wavelength we simply divide the wave's frequency into the speed of sound in air. The formula is $\lambda = V/F$, where λ is the wavelength in feet, V (velocity) is the speed of sound in feet per second, and F is the frequency in Hz. A sinewave with a frequency of 100Hz has a wavelength of 11.3' (1130 divided by 100). At a very high audio frequency of 10,000Hz (10kHz) the wavelength is about an inch.

Notice that a sound's wavelength is dependent on the speed of sound in the medium. For most applications, the speed of sound in air (1130 feet per second) is all we're concerned about. Because sound travels about five times faster in water than in air, the wavelengths of sounds transmitted through water are five times longer.

Phase

The term *phase* describes the fraction of a period that has elapsed in a sinewave. Phase is expressed in degrees, with one complete cycle of a sinewave represented by 360° (see Fig.A-2). Although phase can be measured from any arbitrary point along the sinewave, 0° is usually where the wave begins its positive-going zero crossing.

The term *phase shift* describes a time displacement of the sinewave relative to a fixed point. For example, if the sinewave in Fig.A-2 were delayed by half its period, we would say it experienced a 180° phase shift (because one cycle is 360° , half a period is 180°).

Phase can also describe the time relationship between two or more sinewaves. If one wave is delayed by half a period in relation to a second wave, we say the second wave lags the first wave by 180° . If one wave is delayed by one whole period (360°) in relation to another wave, it as though the phase shift never occurred—the wave is returned to its original phase relationship with the first wave. This is analogous to a clock's second hand returning to the same location every minute.

Phase shift has important consequences when waves are combined. If two waves of equal frequency and the same phase are superimposed, their combination will create a wave of greater amplitude. Specifically, two waves of equal amplitude, equal frequency, and the same phase will yield an increase of 6 decibels (dB) when combined. (Decibels are explained later in this appendix.) These waves are said to *reinforce* each other. If, however, one of these waves is phase-shifted by 180° , the two waves will *cancel* each other and produce no signal. When one wave is at peak positive, the wave shifted by 180° is at peak negative, producing the cancellation.




Fig.A-2 The phase of a wave is its time relationship to a reference point. (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988)

This is exactly what happens in a hi-fi system if *one* loudspeaker is connected with the loudspeaker cables reversed (the red amplifier terminal is connected to the loudspeaker's black terminal, and the black amplifier terminal is connected to the loudspeaker's red terminal). Reversing the leads to one loudspeaker "flips" that signal over, making positive into negative and negative into positive. This is similar to a phase shift of 180°. Now, when one woofer cone pushes forward to create a compression, the second loudspeaker's woofer pulls back to make a

rarefaction. When the compressions and rarefactions are combined, they cancel; the compression from one loudspeaker “fills in” the rarefaction from the second loudspeaker, and less sound is heard. Because the two waves don’t perfectly coincide, they don’t completely cancel, but they do greatly reduce the acoustic output in the bass. If these two signals were combined electrically, however, complete cancellation would occur. Note that if you reverse the red and black leads on *both* loudspeakers, there is no phase shift between the left and right signals, and thus no cancellation occurs.

Absolute Polarity

Some audio products (primarily preamplifiers and DACs) have switches marked “polarity,” “phase,” or “180°.” These invert the polarities of *both* audio channels. The switch “flips” the signals over, making positive into negative and negative into positive. The switch inverts the stereo signal’s *absolute polarity*.

Absolute polarity shouldn’t be confused with *phase reversal between* channels (described earlier). In that case, there was a 180° phase shift in one channel with respect to the other—with severe audible consequences.

Absolute polarity, however, describes a polarity reversal of *both* channels. If musical waveforms were completely symmetrical—that is, identical in the positive- and negative-going halves—absolute polarity would make no difference. You could flip the signal upside-down and never be able to detect the difference, either by listening or by measuring.

But most musical signals aren’t perfectly symmetrical, and are degraded if reproduced with inverted absolute polarity. When a recording is played back with correct absolute polarity, a compression from the musical instrument or voice at the original acoustic event causes a compression from the loudspeaker; a rarefaction from the instrument causes a rarefaction from the loudspeaker. Conversely, playing a recording with inverted polarity results in a compression from the instrument or voice to be reproduced as a rarefaction from the loudspeaker.

It may seem obvious and intuitive that all recordings and playback systems should maintain correct absolute polarity, but that isn’t the case. An audio signal undergoes many reversals of absolute polarity during the recording and reproduction chain—reversals that no one keeps track of. Some recordings have correct absolute polarity and others don’t—there’s about a 50-50 chance that it’s right. Similarly, some playback systems invert absolute polarity and others don’t. The only thing that matters is that the sum total of all the polarity reversals results in correct polarity at the loudspeaker output. If a recording with inverted absolute polarity is played on a system that inverts absolute polarity, the result is correct polarity at the loudspeaker end of the playback chain.

The only way to ensure correct absolute polarity for each recording is by throwing the polarity switch and listening for which position sounds more natural. You can mark your records, CDs, and digital files as to their polarity, and put the polarity switch in the appropriate position for each recording. Reversing

the red and black leads on the cables going to *both* loudspeaker will also invert absolute polarity.

Note that it doesn't do any good to have a non-inverting playback system: roughly half your records and CDs will be reproduced with incorrect polarity whether your system is inverting or not. Moreover, many recordings have no single correct polarity; the disc often contains a mix of inverted and non-inverted tracks. Finally, music recorded with multi-track techniques often contains a mix of inverted and non-inverted signals in the same piece of music, making moot the entire issue of absolute polarity.

Some listeners report radical degradation of the sound when absolute polarity is reversed; others never notice the difference. The audibility of absolute polarity is highly variable, depending on the listener's polarity sensitivity, the instruments on the recording (some instruments produce a less symmetrical wave shape than others), whether or not the recording has a mix of inverted and non-inverted signals (from different microphones), and the phase coherence of the loudspeakers. Loudspeakers with poor time-domain performance can obscure the difference between correct and incorrect polarity.

Complex Waves

Real sounds and musical signals are made up not of sinewaves, but of *complex waves*. Complex waves are many sinewaves of different frequencies superimposed on one another. A vibrating object will produce a *fundamental frequency* (its lowest frequency) overlaid with *overtones* or *harmonics* that are mathematically related to the fundamental frequency. For example, a guitar string vibrating at 440Hz (middle A) will also vibrate at multiples of 440Hz. The second harmonic has a frequency of 880Hz, the third harmonic is at 1320Hz, the fourth at 1760Hz, and so forth.

It is the presence of these harmonics that give an instrument its *timbre* or tonal color. If a piano and a violin both play middle A, both instruments produce the same fundamental frequency of 440Hz. They sound very different because they have different *harmonic structures*: the violin produces stronger harmonics that extend higher in frequency than the piano's.

Harmonics can be *even-order* or *odd-order*. Even-order harmonics are the second, fourth, sixth, eighth, and so on. Odd-order harmonics are the third, fifth, seventh, ninth, etc. The phase relationships of these harmonics and the fundamental have a great influence on the resulting complex wave. The top of [Fig.A-3](#) is a complex wave composed of a fundamental, a third harmonic, and a fifth harmonic, all shown below the fundamental. Just by changing the phase relationship between the fundamental and its harmonics, we get a very different-looking complex wave (top of [Fig.A-4](#)). In this case, the fundamental was shifted by 90° (one quarter of a period) in relation to the third and fifth harmonics.




Fig.A-3 A complex wave is composed of multiple sinewaves (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988).




Fig.A-4 Changing the phase relationship between the constituent sinewaves greatly changes the complex wave's shape. (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988).

A *squarewave* is a sinewave with the addition of the entire series of *odd-order harmonics*. The shape of the complex wave in Fig.A-3 looks somewhat like a squarewave, even though only the third and fifth harmonics are present. If we added seventh harmonic the wave would look even more like a squarewave.

Comb Filtering

A waveform can also be altered if it is combined with a slightly delayed version of itself. This situation occurs when a loudspeaker is placed near a sidewall of a listening room: the listener receives the direct sound from the loudspeaker, and a reflection from the sidewall. The additional distance the reflected sound travels causes the delay. When these two sounds combine and interact, they reinforce at some frequencies and cancel at others. This phenomenon, called *comb filtering*, is a sequence of peaks and notches—hence its similarity to a comb—in the frequency response caused by constructive and destructive interference between the direct and reflected sounds. [Chapter 14](#) includes a detailed description of comb filtering's effects on reproduced sound.

Absorption, Reflection, and Diffusion

When sound strikes a surface, it is absorbed, reflected, or diffused by that surface—or some combination of all three. A soft surface such as carpet or drapes absorbs energy because sound can enter its porous surface to be dissipated within the material's fibers. In this mechanism, sound energy is converted into a minute amount of heat within the absorbing material.

Reflective surfaces are hard and don't allow sound to penetrate the surface; nearly all the energy is reflected back toward the source. Sound is reflected at the same angle at which it strikes; i.e., "the angle of incidence equals the angle of reflection." Light and other electromagnetic waves behave similarly.

Diffusion is a scattering of sound. A diffuser converts a sound traveling in one direction into many lower-amplitude sounds moving in many directions. An irregularly shaped surface tends to diffuse sound, provided that the irregular structures are similar in size to the wavelength of sound striking them. A convex surface acts as a diffuser, but a concave surface does the opposite, concentrating sound into a point.

Diffraction

Diffraction is the bending of sound waves around an object and the re-radiation of energy through an opening. Diffraction makes it possible to hear sounds around corners and behind walls, or through a small opening in an obstruction.

The way sound behaves when it encounters an obstruction, or an opening in an obstruction, changes according to the obstruction's or opening's size in relation to the sound's wavelength. For example, low frequencies encountering an obstruction such as a large wall will bend around that wall, while high frequencies will not. If the obstruction is small compared to the impinging wavelength, it is almost as though the obstruction isn't there; sound just bends around it. The same-sized obstruction would, however, be a barrier to high frequencies.

Diffraction is easily demonstrated by standing behind your loudspeakers; you'll hear bass but no treble. This is because low frequencies can diffract around the loudspeaker cabinet, but mid and high frequencies cannot. The loudspeaker cabinet is small compared to the bass wavelengths, but large in relation to midrange and treble wavelengths.

This phenomenon is shown in [Fig.A-5](#). The following excerpt is the caption accompanying [Fig.A-5](#), reprinted from the McGraw-Hill *Audio Engineering Handbook* by K. Benson, with permission:

“Stylized illustration of the diffraction of sound waves passing through openings and around obstacles. The sound waves are depicted by their wavefronts. The distance between successive wavefronts is the wavelength. The sound waves approaching the barriers are shown as plane waves, as is approximately the case when the source of sound is far away. (a) The wavelength is large compared with the size of the opening and the obstacle. The sound radiates from the small opening as a spherical wave, dispersing equally in all directions and displaying no shadowing behind the small obstacle. (b) The wavelength is small compared with the size of the opening and the obstacle. The plane wave continues through the opening with only slight spreading, and there is a substantial shadow zone behind the obstacle. Although it is not shown, some of the sound would also be reflected back toward the source.”

Diffraction can cause problems in loudspeaker performance: When sound from the drivers reaches the cabinet edge, the energy is re-radiated. This makes it appear acoustically as though the sound is coming from more than one source

(the driver and the cabinet edges). Because this energy re-radiated from the cabinet edges is delayed in relation to the driver's output (it travels slightly farther), the two signals combine to produce peaks and dips in the loudspeaker's frequency response due to constructive and destructive interference (similar to the comb-filtering effects described earlier). Diffraction can be reduced in loudspeakers by rounded baffles, recessed driver mounting bolts, low-profile grille frames, and absorbent material around the drivers. Diffraction is one reason why loudspeakers sound better with their grilles removed.

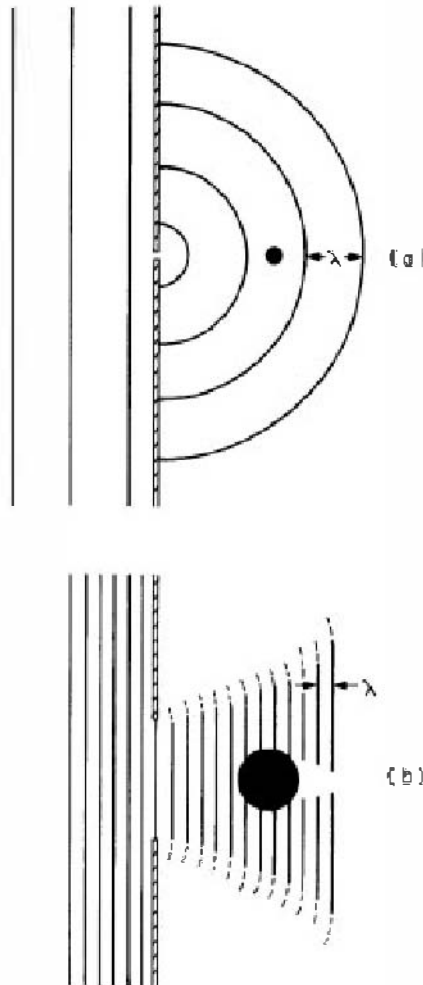


Fig.A-5 How sound behaves when it encounters an obstacle or opening depends on the obstacle's or opening's size in relation to the wavelength of sound. (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988)

The Decibel (dB)

An understanding of the decibel is essential to understanding many aspects of sound and audio. Decibel notation is a convenient way of expressing very large numbers, and also corresponds closely with our perception of sound. A figure in decibels can describe the magnitude of voltages, electrical powers, acoustic powers, and sound pressure levels.

The decibel doesn't express an absolute value, but a *ratio* between *two* values. When we say that one DAC has 2.5dB higher output voltage than another, we are expressing the *difference* in output voltages *between* the players. In the case of our DACs, the decibel notation expresses the *ratio* of their respective output voltages. Any time a decibel notation is used, it must be accompanied by a reference level either to another level, as with the two CD players, or against a standard reference.

This is where the dB notation's suffix comes into play. The letters after the number of dB (80dB SPL, or -20dBFS, for examples) denote the reference level to which the expressed value is referenced. If we say a certain *sound pressure level* (SPL) is 80dB SPL, "SPL" tells us that the sound pressure is referenced to 0dB SPL, the threshold of hearing. By saying "80dB SPL," we mean that the sound is 80dB greater in magnitude than the threshold of hearing. Note that 0dB SPL isn't absolute silence, but the softest sound an average person can hear in a very quiet environment.

The decibel scale is logarithmic, meaning that as the number of dB increases arithmetically, the value that number expresses increases exponentially. For example, each doubling of sound pressure level represents an increase of 6dB, as shown in Table 1.

Sound or electrical power ratio	Decibels	Sound pressure, voltage, or current ratio	Decibels
1	0	1	0
2	3.0	2	6.0
3	4.8	3	9.5
4	6.0	4	12.0
5	7.0	5	14.0
6	7.8	6	15.6
7	8.5	7	16.9
8	9.0	8	18.1
9	9.5	9	19.1
10	10.0	10	20.0
100	20.0	100	40.0
1,000	30.0	1,000	60.0
10,000	40.0	10,000	80.0
100,000	50.0	100,000	100.0
1,000,000	60.0	1,000,000	120.0

Let's say you were in a very quiet room and measured the sound of a fly buzzing at 40dB SPL. We can see in Table 1 that 40dB represents a pressure ratio of 100; the sound of the buzzing fly produces a pressure 100 times greater than a sound at 0dB SPL, the threshold of hearing. The threshold of hearing, represented by the reference level of 0dB SPL, has been established as a pressure of 0.0002 dynes per square centimeter (dynes/cm²). A dyne is a unit of force. The letters "SPL" after the decibel notation tell us that 0.0002 dynes/cm² is the reference level. Knowing this, we can calculate that the sound a buzzing fly produces is a pressure 100 times that of the pressure at 0dB SPL, or 0.02 dynes/cm² (100 × 0.0002).

Two flies would produce a sound pressure level of 46dB—double the pressure, and an increase of 6dB over the 40dB one fly produces. Because 40dB represents a pressure of 0.02 dynes/cm², 46dB represents a pressure of 0.04 dynes/cm². (For the purposes of this illustration, we'll assume the flies produce the same sound pressure, and that their sounds are perfectly in-phase with each other. In reality, the phase relationship between the two sounds will be random, producing a sound *power* increase of 3dB.)

Now consider another example: Let's say that a jet taking off produces an SPL of 120dB. Looking at Table 1, we see that 120dB represents a ratio of 1,000,000—in other words, the jet creates a pressure on our eardrums a million times greater than the pressure at the threshold of hearing. Multiplying 1,000,000 by the reference pressure ($0.0002 \text{ dynes/cm}^2$), we know that the jet creates a pressure of 200 dynes/cm^2 ($0.0002 \times 1,000,000$).

If we add a second jet taking off, we know that the number of dB SPL will increase from 120dB to 126dB, and also that the pressure will double from 200 dynes/cm^2 to 400 dynes/cm^2 . (Again, we'll assume for illustrative purposes that the jets produce equal loudness and are in-phase.)

Obviously, the sound of the second jet taking off is much greater than that of the second fly. But both increased the sound pressure level by the same *ratio*: 6dB. The same 6dB increase represented an increase of only 0.02 dynes/cm^2 in the case of the buzzing fly, but a whopping 200 dynes/cm^2 with the jet.

This logarithmic nature of the decibel makes handling very large numbers easier. These very large numbers are needed to express the vast range of sound pressure levels we can accommodate, from the threshold of hearing to the threshold of pain. The threshold of pain, 140dB SPL, is a pressure on your eardrum of 2000 dynes/cm^2 , or ten million times the pressure at the threshold of hearing. Rather than saying that the threshold of pain is 10,000,000 times louder than the threshold of hearing, it's easier to represent this value as 140dB.

Moreover, this exponential increase parallels our subjective impression of loudness. If we asked a person listening to music at 60dB SPL to turn up the volume until it was "twice as loud," the person would increase the volume by about ten dB. If the person were listening at 100dB SPL and was asked to turn up the volume until it was twice as loud, he would also increase the volume by 10dB. As we can extrapolate from the fly and jet examples, the increase in pressure from 60dB to 70dB is far less than the pressure increase from 100dB to 110dB. Both increases, however, produce similar subjective increases in loudness.

This phenomenon is analogous to the way we perceive pitch: We hear each doubling of frequency as an octave, yet each ascending octave spans a frequency difference twice that of the next lower octave. The octave span between 20Hz and 40Hz is perceived as the same musical interval as the octave between 10kHz and 20kHz, although the latter is a much larger change in frequency.

We can calculate the number of dB between two values with the following formula: $NdB = 20 \log P1/Pr$ where NdB = the number of dB, $P1$ is the measured

pressure, and P_r is the reference pressure. Expressing this formula verbally, we say, "The number of dB equals twenty times the logarithm of the ratio between the two pressures (or voltages, or currents)." This formula works for voltage, electrical current, or sound pressure level. Just as a doubling in SPL is represented by an increase of 6dB, a doubling of voltage is represented by an increase of 6dB.

Using this formula, we can calculate the difference in dB between two DACs having different output voltages. We'll say DAC A puts out 3.6V when playing a full-scale, 1kHz test tone, and DAC B puts out 2.8V with the same test signal. How much difference in dB is there between the two DACs?

We first find the ratio between the two voltages by dividing 3.6 by 2.8. The ratio is 1.2857. By pressing the "log" key on a calculator with 1.2857 in the display, we get the logarithm of the ratio (0.0109). Then we multiply the result by 20, giving us the answer in dB. DAC A's output voltage is 2.18dB higher than that of DAC B.

Calculating power ratios is slightly different. With acoustic or electrical power, an increase of 3dB represents a doubling of power. This is given by the following formula: $NdB = 10 \log P1/P_r$ where $P1$ is the measured power and P_r is the reference power. If we want to find the difference in power output in dB between two amplifiers, we multiply the logarithm of the ratio of the two values by 10, not by 20 as in voltage, current, and sound pressure.

Let's say one power amplifier has a maximum output power of 138W, and another has a maximum output power of 276W. What is the decibel difference between them? First find the ratio between the two powers by dividing 276 by 138. The answer is 2. Find the logarithm of 2 (0.301) and multiply by 10. The answer is that the second amplifier has 3dB more output power than the first amplifier. Note that a doubling of power yields an increase of 3dB.

Frequency, Loudness, and Equal Loudness Contours

So far, we've looked at only the purely physical properties of sound and how it behaves. But sound is meaningless without our perception of it. *Psychoacoustics* is the study of the interaction between the physical properties of sound and our perceptions of those physical properties.

Humans can perceive a frequency range of sound from about 16Hz to 18kHz. The commonly cited range of 20Hz–20kHz is less accurate but easier to remember and express. In reality, most people can hear tones below 20Hz, but few can hear above 18kHz. Teenagers and young to middle-aged women can hear above 18kHz, but that sensitivity decreases with age. It has been proved that women have a greater ability to hear high frequencies than men, and that they retain this ability longer. Although we can't hear sinewaves higher than 20kHz, we *can* hear the effects of limiting the frequency response to 20kHz. The wider the audio system's bandwidth, the faster it can follow steep transient signals. An audio bandwidth of 20kHz isn't wide enough to accurately follow rapid dynamic changes in the music.

This frequency range is wider than it first appears; the frequency span is huge when considering the wavelength of sound. A 16Hz tone has a wavelength of more than 70'; an 18kHz tone has a wavelength of about 3/4".

One of our most fundamental auditory perceptions is *loudness*, the term that describes the magnitude of an auditory event. Loudness is a *perception* that is influenced by factors other than the physical amplitude of sound. In fact, we perceive loudness very differently depending on the sound's frequency. This was first established by two researchers, Fletcher and Munson, who produced a set of curves called the *equal-loudness contours*, also known as *Fletcher-Munson curves*.




Fig.A-6 The perception of loudness is affected by a sound's frequency and volume. (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988)

The equal-loudness contours (Fig.A-6) show how much difference in sound pressure level is required to make all frequencies have the same perceived loudness as the 1kHz reference tone. The numbers running down the middle above each curve are *phons*, a measure of loudness. At the reference frequency of 1kHz, phons are equal to dB.

For example, at a very quiet level of 30 phons (30dB SPL at 1kHz), a bass tone of 50Hz would have to be reproduced at 60dB SPL to be perceived as having the same loudness as the 30dB SPL 1kHz tone. In other words, our ears are less sensitive at the bass and treble frequency extremes than in the midrange.

This large disparity in sensitivity with frequency is more acute at low sound pressure levels, shown by the steep slope of the curves at the graph bottom, and by the gradual flattening of the curves toward the graph top. As the SPL increases, our ears become nearly equally sensitive at all frequencies. At 110 phons (110dB at 1kHz), a 50Hz tone would need to be reproduced at 120dB to be perceived as

being of equal loudness to the 1kHz tone—quite a difference from the 30 phon example.

Everyone has heard this effect; when you turn up the volume on your hi-fi, you hear more bass and treble. The bass seems to extend lower in frequency, and there's more of it. Conversely, if you turn down the volume to where it's barely audible, you hear only midrange sounds. The "Loudness" button on mid-fi receivers boosts the bass and treble to compensate for the ears' reduced sensitivity to bass and treble at low volume.

Weighting Filters

SPL measurements, and measurements of noise in audio equipment, often invoke a *weighting filter* to remove some bass and treble from the reading. Weighting filters attempt to simulate our reduced sensitivity to bass and treble and produce a measurement that more accurately reflects the perceived sound or noise level. Fig.A-7 shows the three weighting curves used in SPL and noise measurements.

A signal-to-noise specification for a preamplifier, for example, may read “85dB, ‘A’ weighted.” This means the noise was filtered with the curve with the characteristics of curve “A” in Fig.A-7 *before* the preamplifier’s noise level was measured. Consequently, applying the weighting curve makes the signal-to-noise ratio appear higher. Similarly, a weighting curve applied to an SPL measurement makes the measured SPL seem lower than the physical magnitude of the sound.




Fig.A-7 Weighting curves applied to SPL or noise measurements approximate the ear’s varying sensitivity with frequency (Courtesy *Audio Engineering Handbook*, K. Benson, McGraw-Hill, Inc., 1988).

The Dynamic Nature of Music

So far we've talked about sound only in the context of steady, continuous tones. But music is dynamic, constantly changing in level, and is composed of impulsive sounds as well as steady ones. A basic understanding of this dynamic quality is important to learning about music reproduction.

We can divide a dynamic sound into three components: *attack*, *internal dynamics*, and *decay* (Fig.A-8). These three aspects of a sound are called the sound's *envelope*. Attack is the way in which a sound begins. The attack can be sudden (a snare drum, for example) or gradual (the build-up of an organ note in a large room). Decay is the way a sound ends. Internal dynamics are the main portion of the sound between the attack and decay. The envelopes shown in Fig.A-8 could be for a snare drum (left figure) or an organ (right). The snare drum has a very sudden attack; it takes very little time for the sound to reach its highest volume. It also has a very short duration and a rapid decay. By contrast, the organ has a very slow attack, taking a longer time to reach its full amplitude. The internal dynamics are very long, and define the note's sound. Similarly, the decay is slow and gradual.




Fig.A-8 Musical sounds are composed of an attack, internal dynamics, and decay.

A sound characterized by sudden attack and rapid decay is a *transient* signal. *Transient response* refers to an audio component's or system's ability to reproduce the steep attack of transient signals. This type of signal has a high *peak-to-average ratio*, meaning that the peaks are much higher in level than the average level. Conversely, the organ has no peaks, so it has a very low peak-to-average ratio. A sound's peak-to-average ratio greatly influences the perceived volume; the ear tends to integrate the sound so that peaks add very little to the overall perceived volume.

Another way of expressing the dynamics of a musical signal is the term *dynamic range*—the difference between the loudest and softest sounds. Technically, a piece of audio equipment's dynamic range is the difference in dB between the device's noise floor and its maximum undistorted signal level. Subjectively, dynamic range is the contrast between a piece of music's quietest and loudest passages.

Localization

Localization is our ability to identify the spatial position of a sound. This ability to localize sounds derives in part from the fact that we have two ears, and that sounds arriving at each ear differ slightly from one another. The brain uses these differences (and other cues) to construct a three-dimensional representation of where a sound source is located.

The most basic mechanism for localizing a sound source is attenuation of high frequencies by the head. Low frequencies can diffract around the head more easily than high frequencies, causing one ear to receive a signal with a gradual attenuation of the midrange and treble. The brain interprets this high-frequency rolloff at one ear as a cue that the sound originates closer to the other ear.

At low frequencies (below about 700Hz), there is virtually no amplitude difference between the left and right ears for a sound emanating from one side. The sound diffracts around the head and the treble-rolloff localization cue doesn't exist. At these low frequencies, a second mechanism takes over in which the phase difference between the left and right ears is analyzed. Because low frequencies have long wavelengths compared to the size of the head, one ear receives a slightly different phase of the sound wave. This information is another cue as to the location of a sound source.

The head, upper body, and outer ear (called the *pinna*) also modify the acoustic signals before they reach the ear. The *pinna*'s unusual shape reflects sounds from its structures into the ear. These reflections combine with sounds directly striking the eardrum, causing comb-filter effects. The brain processes the resultant cancellations and reinforcements to determine a sound source's direction. The modification of the sound by the human body is called the head-related transfer function. Recordings have been made by putting microphones inside a dummy head; the microphones "hear" the signal as modified by the dummy head. In one interesting experiment, recordings were made by placing microphones inside the ear canals of a variety of human subjects. These recordings were then played back by ear buds placed inside the ear canals of the listeners. This technique in essence replaced the head-related transfer function of one listener with that of another. The listening subjects reported wildly different tonal balances and spatial perceptions when listening to another person's head-related transfer function.

The Listening Environment Diagnostic Recording (LEDR) on the first *Chesky Test CD* (Chesky JD37) is a synthesized percussion sound that appears in one loudspeaker, then slowly moves across to the other loudspeaker. Another LEDR test moves the image from the loudspeaker up to a plane above the loudspeaker. The third LEDR test moves the image in an arc from one loudspeaker to the other. These special signals, used to evaluate a listening room, were created by simulating the frequency and amplitude cues imposed by the pinna. By creating notches in the signal's spectrum at precisely calculated frequencies, the signal simulates the effect of a sound source's movement.

You can conduct a simple experiment to demonstrate the pinna's effect on localization. Close your eyes and have a friend jingle keys about four feet away from you, either at head level, waist level, or knee level. Point to where the sound is coming from. Now repeat the experiment with your hands folding your outer ears forward. This eliminates the pinna cues that aid in localization, making it much more difficult to tell where the sound of the keys is originating.

Other Psychoacoustic Phenomena

An interesting phenomenon occurs when we hear a direct sound and a delayed version of that sound (as we would hear an echo). If the echo occurs within about 30ms (milliseconds) of the first sound, the echo isn't perceived as a discrete sound, but instead slightly modifies the sound's timbre. This first 30ms is called the *fusion zone*, and the phenomenon is called the *Haas Effect*.

Similarly, when exposed to a direct and a delayed sound that originate from different directions, we hear the sound as coming from the direction of the first sound. The ear ignores the second sound. Dolby Surround decoding exploits this phenomenon to increase the apparent channel separation between front and rear channels.

Masking is a psychoacoustic phenomenon that renders one sound inaudible when accompanied by a louder sound. Our inability to hear a conversation in the presence of loud amplified music is an example of masking; the music masks the conversation. Similarly, we hear tape hiss or record-surface noise only during quiet passages, not during most of the music. Masking is particularly effective when the masking sound has a similar frequency to the masked component. This kind of masking is called *simultaneous masking*.

Masking also occurs when two sounds are close together in time, an effect called *temporal masking*. The ear/brain tends to hear only the louder of the two sounds—oddly, within a limited time window, it doesn't matter if the louder sound occurred slightly earlier or slightly later than the softer sound. Although it seems logically impossible that a later sound could mask an earlier sound, this occurs because it takes time for the auditory system to process the signal and send it to the brain. When confronted by a louder sound slightly later in time, the louder sound takes precedence.

This appendix is meant only as a cursory survey of the basics of sound and hearing for the interested audiophile; the subject can be studied in great depth. For those interested in learning more about sound and hearing, I recommend the superb introduction to the subject by Dr. Floyd Toole in *Audio Engineering Handbook*, published by McGraw-Hill, ISBN 0-07-004777-4. Further references are included in Dr. Toole's chapters.

Appendix B: Audio and Electronics Basics

A knowledge of electronics and audio technology isn't necessary to choosing and enjoying a high-end music system. You can just select a system, set it up, and not worry about what's going on inside those mysterious black boxes.

Some audiophiles, however, want at least a basic understanding of how audio products work. They have an inherent fascination with audio equipment and want to broaden their knowledge of high-end audio to include a basic conception of electronics and audio circuits. This Appendix is by no means a comprehensive treatment of the subject—that could fill many books. Instead, it provides an overview of audio technology for the interested audiophile.

Moreover, it bridges the gap between the technical descriptions of products published in some high-end audio magazines and the reader who feels intimidated by all the technical jargon. So next time you read that a particular power amplifier has a transformer with dual secondary windings, three full-wave bridge rectifiers, a discrete, direct-coupled JFET Class-A input and driver stage, with a push-pull output stage that switches to Class B operation at 20W RMS, you'll know what it means.

But before we get to all that, let's take a look at the basic vocabulary of electronics and audio technology.

Voltage, Current, Resistance, and Power

The most basic elements of electronics are voltage, current, and resistance. Let's look first at voltage.

A *voltage* exists between two points when one point has an excess of electrons in relation to the other point. A battery is a good example: the negative terminal has an excess of electrons in relation to the positive terminal. When this happens, we say the battery is *charged*.

Voltage is also called *potential difference*, because it has the *potential* to do work and because voltage is always the *difference* in the electrical charges between two points. This is why we say a voltage exists *across* two points, such as a pair of loudspeaker terminals (when driven by an amplifier), or the positive and negative terminals of a battery. One point is more negative (an excess of electrons) than the other point. The *volt* is the unit of measurement that expresses just how great the difference in electrons between two points is. Another way of thinking of voltage is as electrical pressure, like the pressure in a water faucet.

When an electrical *conductor* such as a wire connects two points that have a voltage across them, *current* flows. Current is the electron charge moving through the conductor. The electrons travel through the conductor from the negative charge to the positive charge. If we put a wire across a battery's terminals, current will flow until both terminals have the same number of electrons; the potential difference no longer exists and the battery is discharged, or "dead." Current from a battery is called *direct current*, or *DC*, because the voltage pushing the electrons produces a steady, non-fluctuating electron flow.

If voltage is analogous to water pressure in a faucet, current flow is analogous to water moving through a hose (the conductor). Just as water in the hose is pushed by the pressure of the water behind it, current flows through a conductor pushed by voltage. Current is measured in *amperes*, or *amps* for short.

Resistance is opposition to current flow. If our hose had a kink in it, the flow of water would be reduced. Similarly, an electrical resistance opposes current flow, reducing the number of electrons flowing through the conductor. The higher the resistance, the less the current flow in the conductor (given the same voltage). The unit of electrical resistance is the *ohm*, represented by the Greek letter Omega (Ω).

Ohm's Law

With that introduction, let's look at a simple circuit involving voltage, current, and resistance. Fig.B-1 is a schematic diagram and pictorial representation of a battery connected to a resistor by a conductor. The dashed arrow shows the direction of current flow *through* the resistor, and the two solid arrows point to the voltage *across* the resistor. The voltage source supplies electrons that flow through the conductor and resistor, returning to the positive battery terminal. The current flow will be maintained as long as the voltage exists across the battery.

The amount of current flow is determined by two things: the resistor's resistance value and the battery's voltage. More resistance will decrease current flow. Conversely, a higher voltage will increase current flow. This relationship is expressed by *Ohm's law*, which states that one volt of charge across one ohm of resistance will push one ampere of current. Mathematically, Ohm's law is $E = I \times R$. E is *electromotive force* (another term for voltage), I is current (for "intensity" of electron flow), and R is resistance. If we know any two values, we can calculate the third using Ohm's law. For example, if we know the voltage and current and want to find the resistance, the formula is $R = E/I$. If we want to find the current from the voltage and resistance, we use $I = E/R$.

Let's put some units on our schematic diagram (Fig.B-2). We'll say the voltage is 10V and the resistance is 100 ohms. How much current flows through the resistor? Using $I = E/R$, we divide 10 by 100, getting 0.1 ampere, also expressed as 100 milliamps (100mA). If the resistance is doubled, the current is halved. Similarly, doubling the voltage doubles the current flow. This relationship between voltage, current, and resistance is perfectly linear.

There's another factor here that particularly concerns us in audio—*power*. Power is an expression of the amount of work done by the voltage and current, measured in watts. Power is the product of voltage and current, expressed by Ohm's law as $P = I \times E$ (power equals the current times the voltage). Power is dissipated as heat in the resistor. A resistor is also called the *load*. Going back to our schematic of Fig.B-2, we can calculate the power dissipated by the resistor. We know that the voltage across the resistor is 10V, and the current through the resistor is 0.1A. Multiplying the voltage by the current tells us that the power dissipated by the resistor is 1 watt (10×0.1). Again, the relationship between power, voltage, current, and resistance is linear.


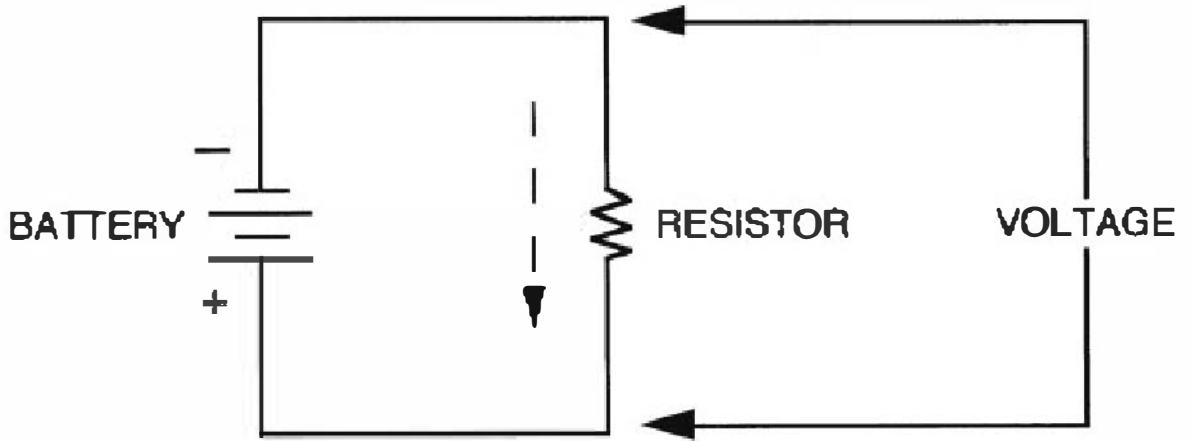


Fig.B-1 A simple series circuit. Electrons from the negative battery terminal flow through the resistor to the positive battery terminal.




Fig.B-2 With an applied voltage of 10V and a resistance of 100 ohms, the current flow through the resistor is 0.1 ampere (100mA).

A power amplifier's output power is a product of its output voltage and current. It describes how much *voltage* the amplifier can develop *across* the load (the loudspeaker input), and how much *current* it can drive *through* the load. Current flow through voice coils in the loudspeaker's drive units makes the cones move back and forth, producing sound.

A perfect amplifier would double its power as the load resistance is halved: it would maintain its voltage and double its current. We know from Ohm's law that if the voltage is kept the same and the current is doubled, the power will double. Such an amplifier is called a *perfect voltage source*. Most amplifiers, however, can't maintain the same voltage across a low-impedance load as they can across a higher impedance, thus the power available to drive a low-impedance loudspeaker is reduced. Amplifiers can also run out of current, again limiting the power delivered to the load.

Series and Parallel Circuits

So far we've discussed a type of circuit called a *series* circuit. In a series circuit, there is only one path through which current can flow. A *parallel* circuit provides multiple paths for current flow. If we add another resistor to our series circuit, we've turned it into a parallel circuit (Fig.B-3). Some of the current flows through the first resistor, some through the second resistor.




Fig.B-3 A parallel circuit provides more than one path for current.

Let's look at series and parallel circuits from the standpoint of a power amplifier driving loudspeakers. An amplifier driving a single loudspeaker is shown in Fig.B-4. The amplifier and loudspeaker form a series circuit; there is only one path for current. If we want to add another loudspeaker to this amplifier channel, we can do it two ways. The first is to put the second loudspeaker in series with the first one (Fig.B-5). The same current flows through both loudspeakers. In a series circuit, the individual resistances are added to find the total resistance. If each loudspeaker presents a resistance of 8 ohms, the total load the amplifier must drive is 16 ohms.




Fig.B-4 An amplifier driving a loudspeaker is a series circuit.




Fig.B-5 Two loudspeakers in series

Another way to connect the second loudspeaker is to wire it in parallel with the first one (Fig.B-6). We've now provided more than one path for current; some current flows through loudspeaker 1, some through loudspeaker 2. Unlike a series circuit, in which the resistances are added to find the total resistance, adding a parallel resistance *lowers* the total resistance (R_T). In the case of our two 8-ohm loudspeakers in parallel, the total resistance the amplifier must drive is 4 ohms. When two equal resistances are in parallel, the total resistance is half that of one resistance. The total resistance (R_T) in a parallel circuit can also be found with the formula $1/R_T = 1/R_1 + 1/R_2 + 1/R_3$, etc.

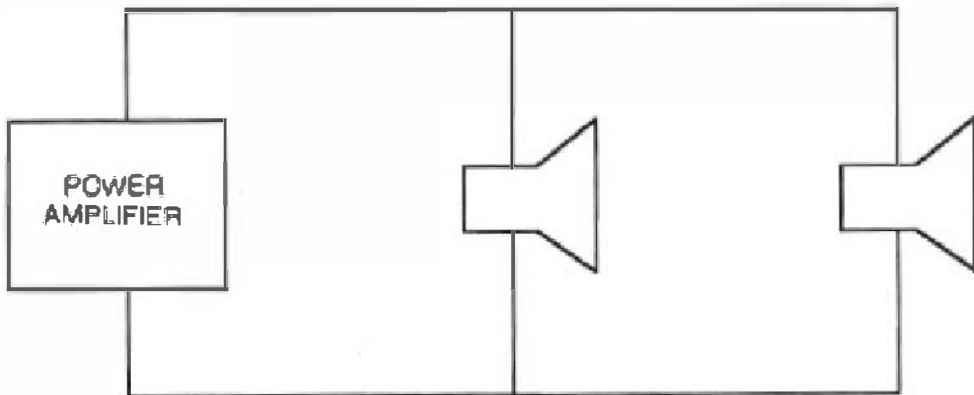


Fig.B-6 Two loudspeakers in parallel

Parallel circuits differ from series circuits in that the voltage across each branch of a parallel circuit is the same as the applied voltage. If we put ten resistors of different resistance values in parallel and apply a source voltage of 20V, that 20V will appear across all ten resistors. However, the current flow will be different through each resistor. It's the current that changes in parallel circuits, not the voltage.

Conversely, in a series circuit, the current remains the same while the voltage across each resistor is different. As current flows through each resistor in series, some of the applied voltage is dropped across each resistor. The amount of that voltage drop is determined by the resistor's resistance value. The sum of the voltage drops across each resistor equals the applied voltage.

Note that, at each resistor in the parallel circuit, the voltage stayed the same while the amount of current flow was different—the opposite of the series circuit.

Alternating Current (AC)

So far we've talked only about circuits using direct current (DC). Another type of current, called *alternating current*, or AC, is the basis for all audio circuits. An audio signal is an example of an AC voltage. An AC voltage reverses polarity, producing current that reverses direction. The principles of DC circuits and Ohm's law described earlier still apply, but AC circuits use additional elements we'll discuss later.

Let's use the AC power line from a wall outlet as an example of AC voltage and current. The 120 volts from the wall socket appear across the socket's two vertical slots. The voltage's polarity (which slot is positive, which negative) reverses itself at the rate of 60 times per second, or 60 Hertz (*Hz*). If we looked at this 120V, 60Hz voltage we'd see the waveform in Fig.B-7.

Starting at the left side of Fig.B-7, the voltage increases until it reaches its positive peak, then decreases to the zero crossing line (no voltage). The voltage then increases again, but in the opposite polarity. After reaching its negative peak, it decreases again. This represents one complete cycle of a sinewave, sixty of which occur every second. The reversing-polarity voltage causes current to reverse direction in whatever conductor is attached to it—usually an AC power cord and power transformer.



Fig.B-7 The 120V, 60Hz AC from a wall socket is a sinewave.

An audio signal is also an AC signal, but unlike our simple 60Hz sinewave, an audio signal is a *complex* AC voltage because it's composed of many frequencies. The AC audio signal is an electrical analog of the sound-pressure wave that created the original audio event. Just as a sound wave is composed of compressions and rarefactions, the audio signal is composed of positive and negative polarity phases.

AC voltage can be expressed in two ways: *peak-to-peak* (p-p) and *root mean square* (RMS). In a p-p expression, the voltage between the waveform peaks is measured. The more commonly used RMS value specifies the voltage at roughly 70% of the peak voltage. RMS better describes the "effective" voltage, or the ability of the voltage to do work. The 120V AC power line has a peak value of

170V, but an RMS value of 120V. When-ever you see an audio signal expressed as a voltage—a preamplifier output level, for example—it is virtually always an RMS value.

Manufacturers of power amplifiers sometimes exploit this difference between peak and RMS to overstate a product's power output. An amplifier can deliver more power on peaks than continuously. [Chapter 5](#) (“Power Amplifiers”) has a more complete discussion of power-amplifier output ratings.

Electromagnetic Induction, Inductance, and Capacitance

We've learned how voltage is equivalent to electrical pressure, current is the flow of electrons "pushed" by the voltage, and resistance is the opposition to current flow. Now let's look at three more electrical properties very important in audio: electromagnetic induction, inductance, and capacitance.

Electromagnetic induction is the relationship between moving electrons and magnetism; moving electrons produce a magnetic field, and a moving magnetic field causes electrons to flow through a conductor. Many audio devices, from phono cartridges to loudspeakers, rely on electromagnetic induction to work.

Let's first look at how a moving magnetic field causes electrons to flow. If you have a conductor, a magnetic field, and relative motion between them, a current (electron flow) will be *induced* in the conductor. For example, if we move a magnet upward through a coil of wire (the conductor), a voltage will be induced across that wire. A pointer on a voltmeter across that coil will swing in one direction when the magnet is moved. When the magnet stops moving, the pointer will return to its original position—no more voltage is being induced across the coil. Now, if we move the magnet in the opposite direction, a voltage will also be induced, but of opposite polarity to that of the voltage induced by moving the magnet up. The voltmeter pointer will now swing to the opposite side of the meter while the magnet is moving. When the magnet stops moving, the pointer will return to its center zero position.

This is exactly how a phono cartridge works. In a moving-coil cartridge, the coils (left and right audio channels) are attached to the cantilever; the cantilever is moved back and forth by the stylus as the stylus follows the modulations in the record groove. The coils sit in a fixed magnetic field. As the coil moves back and forth, a voltage is induced across the coil—the audio signal. The coil cuts through *magnetic lines of flux* surrounding the magnets, inducing a voltage across the coil. The faster the coil cuts these magnetic lines of flux, the higher the induced voltage. The more turns on the coil, the more lines of flux are cut, and the higher the induced voltage. High-output moving-coil cartridges have more turns on their windings than low-output moving-coils. A moving-magnet cartridge works on exactly the same principle, but with the roles reversed: the magnets move and the coils are fixed.

There's another interesting phenomenon of electromagnetic induction: electrons flowing through a conductor create a magnetic field around that conductor. A conductor connected across a battery's terminals will be surrounded by a magnetic field. If we make current flow in the opposite direction by reversing the battery terminals, the north-south magnetic field will also reverse polarity. The magnetic field's density can be increased by coiling the wire to concentrate the magnetic lines of flux, creating a device called an *inductor*, which creates a property called *inductance*. Inductance is measured in henrys (the unit of inductance) and is represented by the letter L. (The schematic symbol for an inductor is shown in Fig.B-8).



Fig.B-8 The schematic symbol for an inductor.

A loudspeaker works on this principle. When current flows through a loudspeaker's voice coil, a magnetic field is set up around that coil. This magnetic field reacts with the fixed magnetic field generated by the magnets surrounding the voice coil. One polarity of the signal-generated magnetic field makes the loudspeaker cone move forward; the opposite polarity makes it move backward. When we apply an AC audio signal to the loudspeaker, AC current flow through the voice coil causes a reversing magnetic field to be set up around the voice coil, making the voice coil move back and forth, and, with it, the loudspeaker's cone.

Let's look more closely at what happens when alternating current flows through a coil of wire (an inductor). The AC current flow creates a magnetic field around the coil that expands, collapses, expands in the opposite polarity, and collapses again. The constantly changing magnetic field follows the AC sinewave voltage applied across the coil. (At DC, there is no expanding and collapsing magnetic field, and thus no induced voltage and no opposition to current flow—the inductor looks electrically just like a simple piece of wire.)

We've just satisfied the requirements for electromagnetic induction: a conductor, a magnetic field, and relative motion between them. The relative motion results from the expanding and contracting magnetic field. This fluctuating magnetic field causes magnetic lines of flux to move through the coil, inducing a voltage across that coil.

This induced voltage opposes the applied voltage, reducing current flow. The higher the alternating current frequency, the greater the induced voltage and the greater the opposition to the applied voltage. At a very low frequency, say 5Hz, the magnetic field expands and collapses very slowly, inducing only a very low voltage across the coil to oppose the drive signal. Conversely, at very high frequencies, the rapidly expanding and collapsing magnetic field induces a much higher voltage across the coil to oppose the drive signal.

This opposition to AC current flow in a coil is called *inductive reactance*. The term *reactance* describes how the coil *reacts* to the applied voltage: the higher the frequency of the applied voltage, the greater the opposition to current flow (the higher the reactance). An inductor can be thought of as a frequency-dependent resistor. The higher the frequency, the more it resists current flow.

Inductive reactance is represented by the term X_L . Inductive reactance is proportional to frequency, expressed in the formula $X_L = 2\pi fL$, where f is frequency in Hz and L is inductance in henrys. For example, a 2-millihenry (2mH) inductor has an inductive reactance of 125.7 ohms at 10kHz. At 20kHz, the inductor's reactance is 251.3 ohms. The higher the frequency, the greater the reactance to the applied voltage.

A power transformer dramatically illustrates how inductive reactance reduces AC current flow. A power transformer consists of two coils of wire placed next to each other, with one coil connected to a 120V AC power line. The DC resistance of that coil is probably about an ohm. If we put 120V DC across one ohm of resistance, Ohm's law tells us that a huge amount of current will flow (120A), which would instantly burn up the transformer. But because that applied voltage is AC, the coil's inductive reactance opposes current flow and prevents the transformer coil from burning. At the 60Hz AC line frequency, the transformer looks like an impedance higher than one ohm.

Another reactive element is called capacitance, represented by the letter C and shown schematically in [Fig. B-9](#). Just as an inductance produces a reactance to an AC signal, capacitance also reacts to AC. While an inductor becomes more reactive at high frequencies, the capacitor becomes more reactive at low frequencies. The capacitor acts as a resistor that increases its resistance as frequency decreases, exactly the opposite of an inductor. A few paragraphs ago we saw that inductive reactance could be calculated with the formula $X_L = 2\pi fL$. To calculate capacitive reactance (X_C), we use the formula $X_C = 1/2\pi fC$. Note that capacitive reactance is a reciprocal function; as frequency increases, capacitive

reactance decreases. At DC the capacitor has infinite reactance and thus looks like an open circuit.



Fig.B-9 The schematic symbol for a capacitor

Filters

We can exploit inductive and capacitive reactance to perform specialized functions in audio circuits. In a two-way loudspeaker, for example, we want only low frequencies to reach the woofer and only high frequencies to reach the tweeter. Consequently, we may design a circuit that looks something like Fig.B-10.

At low frequencies, the capacitor's high reactance opposes current flow. As the frequency increases, the reactance decreases and the capacitor begins to let current flow to the tweeter. At the highest frequencies, the capacitor virtually disappears electrically, allowing maximum current flow. The tweeter is thus protected from low frequencies while high frequencies are allowed to pass. Conversely, the inductor in series with the woofer has a very high reactance at high frequencies, and virtually no reactance at low frequencies.




Fig.B-10 Inductors and capacitors filter the audio signal, as in this loudspeaker crossover.

The capacitor forms a *high-pass filter* and the inductor is a *low-pass filter*. We've just created a simple loudspeaker crossover. (Practical loudspeaker crossovers are more complex.)

A single capacitor or inductor is called a *first-order filter*, producing a *rolloff* of 6dB per octave. Because inductive reactance doubles as the frequency is doubled, the passed voltage is halved when the signal frequency is increased by one octave. Similarly, capacitive reactance is doubled as the frequency is halved by one octave. Because a halving of voltage represents a decrease of 6dB, first-order filters have a rolloff of 6dB per octave. Another term for rolloff is *slope*. First-order filters are also called *single-pole* filters.

Steeper rolloffs can be realized with more inductors and capacitors. *Second-order filters* have a rolloff of 12dB per octave, *third-order filters* roll off at

18dB/octave, and *fourth-order filters* exhibit a 24dB/octave rolloff. Virtually all loudspeaker crossovers have slopes between first-order and fourth-order.

In another combination of a capacitor and an inductor, a *bandpass* filter is created. The bandpass filter will roll off frequencies on either side of the *passband*. A three-way loudspeaker crossover will have a high-pass filter on the tweeter, a low-pass filter on the woofer, and a bandpass filter on the midrange driver.

A filter's *cutoff frequency* is the frequency at which the amplitude is reduced by 3dB. Similarly, the *bandwidth* of an audio device or filter is the frequency span between the -3dB points. All the filters described here are called *passive filters* because they aren't built around an amplifier. Filters that use amplifying devices are called *active filters*.

Capacitors are also used to block direct current, yet pass alternating current. As we'll see later in this chapter, many audio circuits mix DC with the AC audio signal. Putting a capacitor in series with the signal blocks the unwanted DC component and allows the desired AC audio signal to pass. The capacitor looks like an infinitely high resistance at DC, but a low-value resistor to AC signals.

Capacitors can also store a voltage. In fact, the definition of capacitance is the ability to store a charge. If we put a battery across a capacitor, electrons flow from the battery's negative terminal to the capacitor's plate. The voltage source redistributes electrons from one plate of the capacitor to the other. The applied voltage stays across the capacitor even after the battery is removed.

This stored charge can be lethal. A power amplifier with large power-supply capacitors can kill you even if the amplifier hasn't been plugged in for years. Even small capacitors can give you a shock. The storage ability of capacitors can be demonstrated by turning off your power amplifier while it is reproducing music. It takes several seconds for the music to fade after the amplifier's power source has been removed. The amplifier is running off the charge stored in its capacitors.

Impedance

Impedance is the opposition to current flow in an AC circuit, specified in ohms and represented by the symbol Z . Impedance is to an AC signal what resistance is to DC. If we're dealing with DC, opposition to current flow is resistance. If the voltage is AC, opposition to current flow is impedance. Because loudspeakers are driven by AC signals, they have an impedance, not a resistance.

Impedance differs from resistance in that impedance implies that the load is not a simple resistance, but a combination of resistance, inductive reactance, and capacitive reactance.

An audio playback chain is a series of impedances. Every component with an input presents an impedance to the signal driving it. The impedance a component presents to the drive signal is called its *input impedance* (also called *load impedance*). A DAC's analog output drives the preamplifier's input impedance, the preamplifier drives the power amplifier's input impedance, and the power amplifier drives the loudspeaker's impedance. Impedance can be thought of as the load through which the component must force current.

Components in an audio chain also have an *output impedance*, or *source impedance*. If you think of a component—a CD player, for example—as an audio signal generator, its output impedance is like a resistance in series with that signal generator.

Generally, audio products have output impedances of between a few tens of ohms and a few hundred ohms. Input impedances, however, are very high, usually between 10,000 ohms (10k ohms) and 1,000,000 ohms (one Megohm, or one million ohms). A product with a 1M ohm impedance will draw less current through it than one with a 10k ohms input impedance. The higher the input impedance, the easier it is for the source component to drive; the input impedance is less of a load on the source component.

Let's combine our knowledge of Ohm's law, capacitance, and impedance to examine the interaction between a preamplifier and a power amplifier through a long run of interconnect. Fig.B-11 shows this circuit pictorially and schematically. Because all cables and interconnects have some inherent capacitance, we can represent the interconnect capacitance by the symbol for a capacitor between the signal conductor and ground. The preamp's output impedance is represented by the resistor in series with the signal generator. The power amplifier's input impedance is represented by the resistor inside the power amplifier outline.

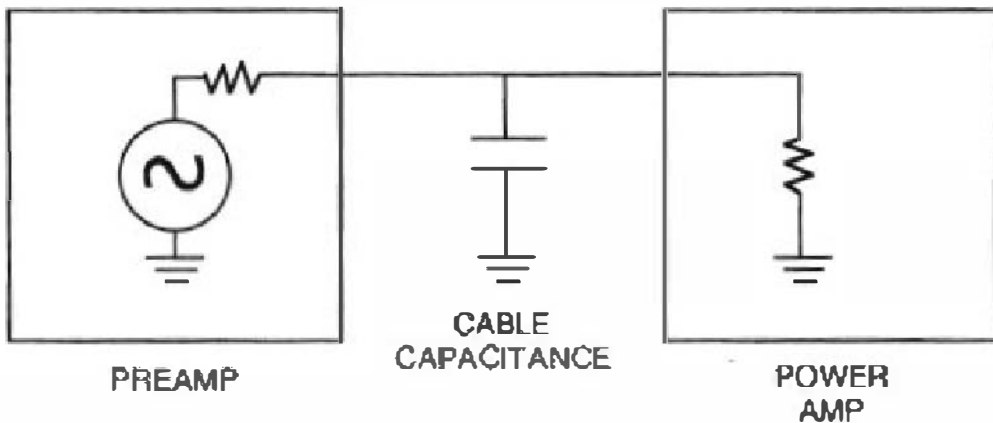



Fig.B-11 A preamplifier-to-power-amplifier connection shown pictorially and schematically. The interconnect's capacitance is shown as the capacitor in parallel with the power amplifier's input.

Let's say that the preamp has a very low output impedance, the power amplifier has a high input impedance, and the cable capacitance is low. All these are ideal conditions. The power amplifier's high input impedance means very little current will flow, causing very little voltage drop across the preamplifier's output impedance. Because the cable's capacitance is low, its reactance will remain high up to a very high frequency. Everything works just fine.

But if the preamplifier has a high output impedance and the cable is excessively capacitive, we run into problems. At low frequencies, the cable capacitance has no effect: its capacitive reactance is extremely high. In fact, it's almost as if the capacitor weren't there. But as frequency rises and the capacitive reactance drops, the capacitor looks like a resistor connected to ground. The higher the frequency, the lower the value of the "resistor" connected between the signal and ground. Consequently, as the signal frequency increases, more and more of the drive signal is dropped across the source impedance and less and less across the load (the power amplifier's input). We've just created a low-pass filter.

The result is high-frequency rolloff across the power amplifier's input terminals. The capacitor provides high frequencies with an easier path to ground than does the amplifier's input circuit.

Moreover, the greater current flow caused by the power amplifier's low input impedance and the cable capacitance causes the preamplifier's output stage to work harder, perhaps even becoming *current-limited*. This condition occurs when the source can no longer supply the current through the load, resulting in a clipped waveform.

The result of an impedance mismatch and high interconnect capacitance between a preamplifier and power amplifier is a rolled-off treble, lack of air and extension, mushy bass, and loss of dynamics.

We can calculate the -3dB rolloff point for our low-pass filter using the formula $F = 1/2 \pi RC$. F is the -3dB point, R is the power amplifier's input impedance in parallel with the source impedance, and C is the interconnect capacitance in farads. Interconnect capacitance is measured in picofarads (10^{-12} farads) per foot, with less than 100pF per foot considered low capacitance.

This example illustrates what can happen with an impedance mismatch between components, particularly if the cable is capacitive. Source components should have an output impedance of less than a few hundred ohms, preferably a few tens of ohms. Input impedances should be at least 47k ohms (a standard value).

Loudspeaker Phase Angle

Another application of our knowledge of impedance, capacitance, and inductance is understanding the concept of *phase angle* in a loudspeaker's impedance. Phase angle describes how reactive a load is compared to the load's resistance.

A loudspeaker impedance measurement includes the impedance magnitude plotted in ohms, but also a separate plot of how reactive the loudspeaker is to the drive signal as a function of frequency. This reactance is caused by the inductors and capacitors in the loudspeaker's crossover. Fig.B-12 is a loudspeaker's impedance magnitude and phase-angle plot. If we performed this measurement on a resistor (which exhibits no inductive or capacitive reactance), we would see straight lines on the impedance-magnitude and phase-angle traces (the latter is dotted). A pure resistor has no phase angle. But because the loudspeaker is reactive, its phase angle and impedance change with frequency.

Phase angle is introduced because capacitive reactance causes voltage to lag slightly behind the current flow in time. In an inductor, the opposite occurs: current lags voltage. Capacitive reactance creates a phase angle of -90° , while inductive reactance creates a phase angle of $+90^\circ$. When capacitive and inductive reactance are combined, the phase angle may be anywhere between $+90^\circ$ or -90° . When the dotted line on a phase angle plot is above the 0° center line, the loudspeaker is primarily inductive. When the dotted line dips below the 0° center line, the loudspeaker is primarily capacitive. The less severe the phase-angle swings (deviations from a flat line), the less reactive the loudspeaker and the easier it is for the amplifier to drive. A phase-angle trace that stays near the 0° center line makes the loudspeaker look more like a simple resistor to the power amplifier.

Impedance Magnitude (2 ohms/vertical div.) & Phase (22.5 degrees/vertical div.)




Fig.B-12 A loudspeaker's impedance magnitude and phase angle. (Courtesy *Stereophile*)

Capacitor Types

Capacitors are crucial elements in audio circuit design and manufacture. They are used everywhere in audio products, and can significantly influence the product's sound. Capacitors vary greatly in their physical and electrical properties; some are better than others for use in audio circuits. There are also huge price differences between good and poor capacitors. Some capacitors found in the best loudspeakers can cost as much as \$400 each.

Capacitors are formed of two plates separated by an insulator called the *dielectric*. The dielectric can be air, Mylar, ceramic, paper, oil, polystyrene, polypropylene, foil, Teflon, tantalum, or other materials. Capacitors are referred to by their dielectric material: we may say a particular preamplifier uses Teflon capacitors throughout. The best audio capacitors are Teflon, gold foil, silver foil, polystyrene, and polypropylene.

One type of capacitor found in power supplies is the *electrolytic* capacitor. These are usually large metal cans that have very high capacitance values. Electrolytics work well in power supplies, but are poor choices for capacitors in the signal path. One manufacturer even gets rid of all electrolytics from the power supply, using much more expensive polypropylene caps instead.

Anatomy of an Audio Product

With this basic knowledge of electronics, let's see how these principles are applied in audio products.

The Power Supply

Every audio component that plugs into a wall outlet contains a power supply. The power supply converts 120V, 60Hz AC power from the wall outlet into the lower DC voltages needed by audio circuits. This is a critical job in high-end audio: the power supply has a large influence on the product's sound quality. Often, the power supply accounts for much of a component's size, weight, and cost. The general thinking in high-end product design is to have extremely simple audio circuits and extremely elaborate power supplies. We want a pure and simple path for the audio signal to follow, and a very clean and stable DC supply. The first component in the power supply is the *power transformer*. This device *steps down* the 120V AC powerline to a lower AC voltage, perhaps $\pm 30V$ (we'll see how it does this in a minute). The transformer's schematic symbol is an accurate representation of its construction: it consists of two coils of wire placed next to each other (Fig.B-13).




Fig.B-13 The schematic symbol for a transformer represents its physical structure.


A little earlier we learned that whenever current flows through a conductor, a magnetic field is set up around that conductor. We also learned that whenever we have a conductor, a magnetic field, and relative motion between the two, a voltage is induced. Both principles are in operation in a transformer. When the 60Hz AC flows through the transformer's first coil, called the *primary winding*, an expanding and collapsing magnetic field is set up around this coil. This expanding and collapsing field created around the primary winding induces a voltage across the second coil, called the *secondary winding*. The 60Hz AC signal thus appears across the secondary winding, even though there is no physical connection between this winding and the AC wall outlet.

The voltage induced across the secondary is determined by the *turns ratio* between the primary and secondary. By changing the turns ratio, any value of

output voltage can be obtained. Some transformers have multiple secondary windings to supply different circuits. The transformer thus converts 120V, 60Hz AC to a lower 60Hz AC voltage or voltages. The transformer also *isolates* the component from the power line.


Transformers are rated by power capacity, called a VA rating (Volt/Amperes). A transformer rated at 1kVA (1000VA), for example, can deliver 100V at 10A (assuming a 1:1 turns ratio). Transformers in low-signal components (preamps, CD players) are typically rated at a few tens of VA, while the largest power-amplifier transformers may reach a whopping 5kVA. Such a transformer may weigh more than 100 pounds.

Many audio products use *toroidal* transformers instead of the more conventional *laminated* type. A toroidal transformer is doughnut-shaped to concentrate its magnetic field in a tighter configuration. Toroidals are more expensive than laminated types, but are more efficient and radiate less noise into the surrounding circuitry. Just as the magnetic field surrounding the primary winding induces a voltage across the secondary winding, audio circuitry near the transformer can pick up 60Hz *hum*. Toroidal transformers reduce this risk. The stepped-down voltage from the transformer's secondary winding is still AC, but we need DC to supply the audio circuits. The next element in the power supply is the *rectifier*, a device that turns AC into DC. Rectifiers are made from diodes, the most basic *semiconductor* element. A diode allows current to flow in one direction but not the other. Fig.B-14 shows the effect of putting a single diode on a transformer's output. The diode allows one polarity to pass, but blocks the 60Hz sinewave's other half. (Which half it passes—positive or negative—is determined by the diode's direction.) In effect, the diode removes one polarity of the signal. Because the resultant voltage never changes polarity, it is, technically, direct current. This method, called *half-wave rectification*, isn't very efficient. But by putting four diodes together, we can make a *full-wave bridge rectifier* (Fig.B-15). In effect, the other sinewave polarity is "flipped over." The full-wave rectifier produces smoother DC and delivers more power to the supply.



(A) Transformer feeding a half-wave rectifier

Fig.B-14 A half-wave rectifier passes only one phase of the AC powerline.



(C) Transformer feeding a bridge rectifier

Fig.B-15 A full-wave rectifier “flips over” one polarity of the AC powerline. (Figs. B-14 and B-15 copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen)

The rectifier output is DC (it never changes polarity), but still has AC *ripple* riding on it. Ripple consists of vestiges of the AC signal superimposed on the DC; the DC shifts up and down at 120Hz. Note that a full-wave rectifier produces ripple at 120Hz, not 60Hz, because the negative phase of the sine wave is “flipped over” (shown earlier in Fig.B-15). Ripple is reduced by adding *filter capacitors*,

whose storage function smooths the AC ripple riding on the DC voltage. In fact, filter capacitors are also called *smoothing* capacitors.

The filtered output is now DC with much less ripple, but it still isn't ready to supply our audio circuits. Any changes in the input line voltage will affect the DC voltages supplying the audio circuits—an unacceptable condition. The job of making this DC more stable and at the correct voltage falls on the *voltage regulator*. A voltage regulator maintains a constant output voltage regardless of changes to its input voltage. The voltage regulator's input is DC that may have some ripple and be subject to fluctuation and drift; the regulator's output is a pure, stable, and unwavering DC source. Voltage regulators are fed a voltage slightly higher than the desired output voltage; they can thus maintain a stable output voltage even if the input voltage changes slightly. A typical 12V regulator (the output voltage) may have an input voltage of 15V. This 12V is now good enough to become the power-supply *rail* for the audio circuits. Voltage regulators usually have heatsinks attached to them; the difference between input and output voltage is dissipated as heat.

All these elements—power transformer, rectifier, filter capacitors, and a voltage regulator—are shown schematically in Fig.B-16, and pictorially in Fig.B-17.




Fig.B-16 A simple power-supply schematic

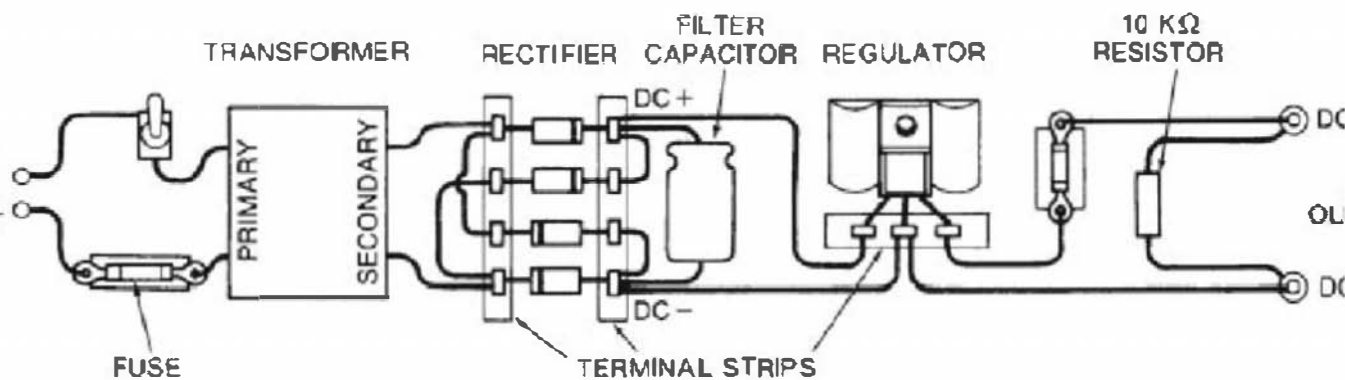


Fig.B-17 The power supply shown pictorially

Figs.B16 and B-17 copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen

The power supply will often produce several DC voltages of different polarities (positive and negative) to supply different parts of the circuit. At the very minimum, a digital processor will need +5V to the digital circuits and $\pm 12V$ to the analog stage. Most processors have at least four *regulation stages*, with some having up to 15 separately regulated supply voltages. Multiple regulation stages help *isolate* circuits from each other. Different audio circuits can interact with each other through a common regulation stage, degrading performance. For example, if a CD player's tracking-servo power supply also fed the analog output stage, the supply voltage would *droop* when the tracking servo demanded current, affecting the analog circuit's performance. Designers agree: the more regulation stages, the better.

Most regulators are integrated circuits (see Fig.B-17). Conversely, some regulation stages are called *discrete* because they use separate transistors, resistors, and capacitors to form the regulator. Discrete regulation is much more expensive to implement, but is generally considered to produce better-sounding products than those using IC regulation. Some very elaborate products use *cascaded* regulation, in which the output of one regulator supplies the input to the second regulator. This ensures a purer, more stable DC supply. The ultimate power supply uses many cascaded discrete regulation stages.

The regulated DC voltages are often *bypassed* with very-low-value ($0.01\mu F$) high-quality capacitors. The capacitor is placed physically next to the supplied circuit, and electrically between the power-supply rail and ground. The capacitor provides a path to ground for noise or power-supply-rail disturbances.

Other techniques that contribute to good power-supply performance are high values of filter capacitance, with some of the capacitors located next to the audio

circuits they supply. A large transformer is also a plus, even though the VA rating may already exceed the circuits' current requirements. Most designers agree that the bigger the transformer, the better. Separate transformers are also an advantage, further isolating individual circuit elements from each other.


Power supplies are a critical element in the music reproduction chain. Although mediocre supplies will still make the product work, the very finest components are distinguished by elaborate and expensive power supplies. Much of what you pay for in high-quality audio products is in the power supply.

Now let's look at some circuits that the power supply feeds—transistor and tube amplifiers.


Amplifier Circuits

Amplifier circuits are the basic building blocks of preamps, power amps, phono stages, and DACs. If you understand a simple, one-transistor amplifier, you'll have a good grasp of how all amplifiers work.

A transistor is a solid-state semiconductor that has three elements: the *base*, *emitter*, and *collector*. It can be thought of as a sandwich of these three layers, shown schematically in Figs.B-18a and b. The transistor converts a small signal at the base to a large signal at the collector. If you think of a hose with water flowing through it, the water source is the emitter, and the base is a kink in the hose that opens and closes to allow more or less water to flow.



(A) NPN transistor



(B) PNP transistor

Figs.B-18 a and b A transistor's schematic symbol (NPN and PNP types). (Copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen)

Let's see how this works by looking at the single-stage transistor amplifier shown in Fig.B-19. The DC power supply is connected to the collector at the top of the figure. The emitter is negative with respect to the collector. Because a voltage exists between the emitter and collector, current wants to flow up through the transistor. The base can prevent or permit current flow, depending on the base voltage. The base is like a control valve that can turn current flow on and off. When the transistor is *turned on* by a tiny base current, a large current flows from the emitter, through the base, and into the collector. When the transistor is *turned off* by a lower base current, no current flows from the emitter to the collector.

If we put an audio signal on the base, current flow from the emitter to the collector will be controlled by that audio signal. When the audio signal is of higher amplitude, the transistor will be more turned on, allowing more current to flow from the emitter to the collector. A low-amplitude audio signal will allow less current to flow through the transistor. A tiny base current thus controls a

large current through the emitter and collector. This is how a transistor amplifies an audio signal.




Fig.B-19 A simple one-transistor amplifier. (Copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen)

The amount of amplification provided by a transistor is called *gain*, expressed either as the ratio between the input and output voltages, or in decibels (dB). A gain of 10 means that the output signal has ten times the voltage of the input signal; a 1kHz sinewave with an amplitude of 10mV applied to the base will produce a 1kHz sinewave at the collector with an amplitude of 100mV.

In practice, the transistor is partially turned on all the time. A small base current, called the *bias current*, is always present. Bias is essential to making the transistor work. Bias makes the transistor conduct *quiescent current* through the emitter and collector. Quiescent current is the current that flows through the transistor's collector when no input signal is applied to the base. The amount of bias current is carefully selected so that the transistor operates in its linear operating region (**Fig.B-20**). Bias allows the transistor to amplify both positive and negative phases of the audio signal; because the supply voltage is positive and the emitter is connected to ground, the transistor can never produce a negative voltage at the collector. By making the transistor conduct all the time with bias, the zero crossing point is some voltage between ground and the supply voltage (determined by the values of the emitter and collector resistors). When the input

signal swings negative, the output voltage is simply “less positive” than this zero crossing voltage.




Fig.B-20 Bias makes the transistor operate in its linear region. (Copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen)

As explained in the section on series circuits, two resistances in series will produce a voltage drop across the resistors. This is what happens in a transistor circuit: current flow through the collector and emitter resistors form a *voltage divider*, producing a DC voltage at the collector. When we apply an AC audio signal on top of the bias current, the transistor becomes more and less turned on, amplifying the base current. The signal at the collector is thus a DC voltage with the audio signal riding on it. The final output is an amplified replica of the signal input to the base. In the configuration shown here, called a *common-emitter* amplifier, the output signal is 180° out of phase with the input signal.

Note that the large signal at the output is simply the DC power supply converted to an AC audio signal by the transistor. This is an important point for audio products: the power supply is effectively in the audio signal path. It's easy to conceptualize the power-supply rails as a static DC supply outside the audio

signal path. But because the audio signal appears at the collector, and the power supply is also connected to the collector, anything that happens to the power supply—noise, high output impedance, fluctuations—also happens to the audio signal. This is why high-end amplifiers have such elaborate power supplies.

In Fig.B-19, R_1 is the load resistor, R_2 is the emitter resistor that stabilizes the transistor's operation, R_3 and R_4 create a voltage divider that applies the correct bias to the base, and C_3 is a bypass capacitor. C_1 and C_2 are *coupling capacitors*, which block DC from getting into the transistor (C_1) and the next amplification stage (C_2). Because capacitors in the signal path are generally not a good thing, some designers prefer to remove them. An amplifier circuit without coupling capacitors is called *direct-coupled*. If no coupling capacitors are used, DC can be prevented from reaching the amplifier's output with an active device called a *DC servo*.

A technique for decreasing distortion in amplifiers is called *negative feedback*, or just *feedback*. A portion of the amplifier's output signal is returned to the input, reducing distortion. When the feedback *loop* is around a single or several gain stages, the feedback is called *local*. When the feedback is taken from the amplifier's final output and returned to the input stage, the feedback is called *global*. Excessive feedback produces lower distortion figures, but often makes the amplifier sound worse. High-end amplifier designers try to use as little feedback as possible.

Going back to Fig.B-19, we can see that the transistor amplifies the entire audio signal; consequently, it's always turned on—i.e., always conducting current. Because the transistor is always turned on, this amplifier operates in *Class-A* mode. Class-A simply means that the transistor amplifies the entire audio signal and is always turned on.

Although it may seem intuitive that transistors would work this way, virtually all power-amplifier output stages use a more efficient technique called *Class-A/B*. Class-A/B operation produces less heat and is more efficient than Class-A, making Class-A/B stages more economical in circuits that must handle lots of current.

In Class A/B, transistors are arranged in *complementary pairs*. One half of the complementary pair is an *NPN* transistor, the other half is a *PNP* transistor. The NPN/PNP designation refers to the transistor's internal structure, which determines the voltage polarities needed to operate the transistor. In a complementary pair, an NPN transistor is connected to a PNP transistor (Fig.B-21). The “top” transistor will handle the positive half of the audio signal,

and the bottom transistor will amplify the negative portion of the signal. When the top transistor is conducting, the bottom transistor is turned off, and vice versa. This technique keeps the transistors turned off half the time, producing less heat and reducing stress on the transistors.

The signal driving a complementary pair is processed by a *phase splitter* before the transistor bases. A phase splitter duplicates the audio signal so that the positive phase drives the NPN transistor of the complementary pair, and the negative phase drives the PNP half of the pair.

A Class-A/B amplifier is also called a “push-pull” output stage because one transistor “pulls” current while the other transistor “pushes.” Current flow is indicated by the directions of the arrows in Fig.B-21.




Fig.B-21 A complementary output pair of transistors driven by a Class-A driver stage. (Copied with permission from *Audio Technology Fundamentals* by Alan A. Cohen)

A true push-pull amplifier actually operates in pure Class-B, but virtually all power amplifiers operate in Class-A up to a few watts, then switch to Class-B operation, hence the “Class-A/B” designation. At low signal levels, both halves of the complementary pair always conduct, meaning that the output stage is operating in Class-A. The power output level at which the amplifier switches into Class-B is determined by the bias current applied to the output transistor bases. More bias current turns the transistors on harder, allowing them to handle both phases of the audio signal up to a higher power output. There’s a limit to how much power the amplifier will deliver in Class-A; the designer just can’t keep increasing the bias current. Maximum Class-A power output is determined by how much current the transistors can handle, how well heat is removed from the transistors, and how much current the power supply can deliver. Pure Class-A amplifiers are huge, very heavy (due to their massive power transformers), generate lots of heat, and are very expensive in terms of “dollars per watt.”

Class-A operation produces less distortion than Class-A/B, and generally sounds better. A kind of distortion in Class-B amplifiers called *crossover distortion* can produce a waveform discontinuity at the zero crossing point where one transistor turns off and the other turns on. The two halves of the complementary pair aren’t biased identically, producing crossover distortion. This condition can’t occur in Class-A amplifiers. Class-A output stages have other another big advantage: the transistors are always at the same temperature regardless of the signal they are amplifying, making the output stage more stable and linear. (Chapter 6 contains a discussion of amplifier output stages and their musical qualities.)

The NPN and PNP transistors described previously are called *bipolar* devices because current flows through both polarities (P and N) of semiconductor material. Another type of transistor, called the *Junction Field Effect Transistor* (JFET), is often used in small-signal applications. The JFET is called a *unipolar* device because current can flow in only one direction. A JFET has a *gate*, a *source*, and a *drain*, which correspond to the base, emitter, and collector in our NPN and PNP transistors. JFETs are often used in amplifier input stages for their high input impedance. JFETs are also turned on (biased) by voltage rather than by current, as are bipolar transistors. JFETs are found only in small-signal circuits. In addition, JFETs are also used as electronic switches.

Another type of field-effect transistor is called the MOSFET (MOS stands for Metal Oxide Semiconductor). MOSFETs are sometimes used in power-amplifier output stages.

Amplifier Distortion

All the amplifier circuits we've talked about produce distortion—their output signals aren't exact replicas of their input signals. Instead, the signals are changed in certain ways. In *harmonic distortion*, harmonics of the input signal appear in the output. For example, if a 100Hz sine-wave is input to the amplifier, the output will be a 100Hz sine wave, but with the harmonics of 100Hz (the second harmonic at 200Hz, third harmonic at 300Hz, fourth at 400Hz, etc.) added to the signal.

Harmonic distortion is expressed as a percentage of the output signal. The term *Total Harmonic Distortion*, or THD, is an expression of the percentage of all the harmonic components combined. For example, if the amplifier puts out 10V and the harmonic distortion products are 10mV, the amplifier has 0.1% THD. Although THD is a common amplifier specification, a more instructive analysis of harmonic distortion involves examining *which* harmonics are produced by the amplifier. Second- and third-harmonic distortion components aren't nearly as objectionable as upper-order harmonic distortion components such as the fifth, seventh, and ninth. In fact, 10% second-harmonic distortion is less musically objectionable than 0.5% seventh-harmonic. We can quickly see that a single THD figure tells us very little about the audibility of an amplifier's distortion. All amplifiers produce different harmonic distortion spectra—one possible reason why they sound different.

Another distortion in amplifiers is called *intermodulation distortion*, or IMD. When a signal of two frequencies is input to an amplifier, the amplifier will output the two frequencies, plus the *sum* and *difference* of those frequencies. For example, if we drive an amplifier with a mix of 100Hz and 10kHz, the amplifier's output will be 100Hz and 10kHz, *plus* 9.9kHz (the *difference* between 100Hz and 10kHz) and 10.1kHz (the *sum* of 100Hz and 10kHz). In fact, intermodulation distortion produces an infinite number of intermodulation products. Like harmonic distortion, intermodulation distortion is expressed as a percentage of the output signal.

Amplifiers can also distort the audio waveform's transient character, a distortion called *slew-rate limiting*. If the steepness of the input signal's attack is faster than the amplifier's ability to respond, slew-rate limiting will occur. An amplifier's slew rate, measured in volts per microsecond ($V/\mu s$), describes how much voltage the amplifier can swing in a specified time period ($1\mu s$). If the amplifier can swing $20V/\mu s$ and the input signal has transient signals that swing

faster than $20\text{V}/\mu\text{s}$, the amplifier will distort the music's dynamic structure. Slew rate is often referred to as an amplifier's "speed."

Vacuum Tube Amplifiers

Vacuum tube amplifiers are very similar in operation to transistor amplifiers, but are generally much simpler. The three elements in a tube are the *control grid*, *cathode*, and *plate*, shown schematically in Fig.B-22. These correspond to the base, emitter, and collector of the bipolar transistor. The cathode emits electrons, the plate collects them, and the control grid controls how many electrons flow through the tube from cathode to plate. This most basic form of tube is called a *triode*.

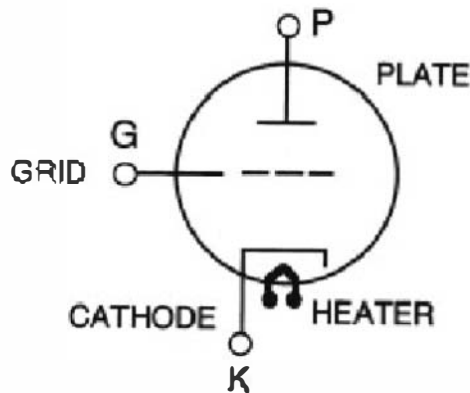


Fig.B-22 The three elements of a vacuum tube are the grid, cathode, and plate.

One more element is needed to make the tube work: the *filament*, or *heater*. The heater is the part of the tube that glows, heating the cathode to a high temperature. When the cathode is heated, electrons are given off, a phenomenon called *thermionic emission*. These electrons “boiling” on the cathode surface want to flow through the vacuum to the positively charged plate, but the grid is in the way. The voltage applied to the grid determines how many electrons jump from the cathode to the plate. The grid is a mesh of wire that produces an electric field that controls the flow of electrons, yet doesn’t impede them on their way to the plate. By applying an audio signal to the grid, we can control a large current flow through the tube.

In some vacuum tubes the cathode is heated directly rather than indirectly by a separate heater. These so-called direct-heated tubes are sometimes used in single-ended triode amplifiers.

In addition to amplifying audio signals, triodes are also used as *out put buffer* amplifiers in a configuration called a *cathode follower*. The cathode follower has a

voltage gain of less than one (meaning it provides no amplification), but has other attributes, such as a high input impedance and low output impedance. In tubed digital processors and preamplifiers, cathode followers are often the last stage before the output jacks.

Another tube type, called a *pentode*, is often used in amplifier output stages. A pentode is essentially a triode with two additional elements, called the *screen grid* and the *suppressor grid*. Both of these elements are between the control grid and plate. When connected to a voltage, the two additional grids affect the electric field within the tube, and provide much higher voltage gain than is possible with a triode. While triode voltage gain may be between 10 and 50, pentodes can provide voltage gains of up to 500. Most low-level tubed circuits—such as preamplifier input, gain, and output stages, and power-amplifier input and driver stages—use triodes for their better sound. Pentode stages are usually reserved for tubed output stages where high amplification is needed. Some tubed power amplifiers use triode output stages, but their power output is relatively low. However, they sound extremely good.

Tubed amplifiers usually have much higher power-supply voltages than do transistor amplifiers. While the collector-supply voltage in a transistor amplifier may be 15–100V, tube plate voltages can run as high as 800V. The plate voltage is often called the “B+” supply. Further, tubed amplifiers need a high current supply—usually 6V or 12V—for the tubes’ heaters. (See [Chapter 6](#) for more discussion of tubed power amplifiers.)

Operational Amplifiers

An *operational amplifier*, or *op-amp*, is an integrated circuit capable of amplifying audio signals. The op-amp replaces many transistors, capacitors, and resistors with a tiny plastic or metal package. Very few external components are needed to build an audio circuit based on op-amps—just supply them with a DC power supply, a *feedback resistor*, and an input signal, and they work.

In theory, op-amps are ideal for audio circuits: they have very high input impedance, very low output impedance, and can provide high gain over a wide bandwidth. They are much easier and cheaper to implement than *discrete* circuits made of transistors, capacitors, and resistors. Unfortunately, most of them don't sound as good as discrete circuits. The very best audio components use discrete circuitry, while budget and moderately priced products usually use op-amps. Some products employing op-amps have sounded superb, however.

An op-amp has two inputs, one output, and pins for the DC power supply. The opamp is typically supplied with $\pm 12\text{V}$ DC from the power supply. It has two input terminals, marked + and –, called the *non-inverting input* and the *inverting input*, respectively. A signal input to the non-inverting input will produce an output signal of the same polarity as the input signal. Conversely, an input signal at the inverting input will produce an output signal of opposite polarity to the input signal. Op-amps can thus function as inverting or non-inverting amplifiers. The value of the feedback resistor determines the op-amp's gain.

Op-amps can be used as voltage amplifiers, active filters, or output buffer amplifiers. (They also have thousands of uses not applicable to audio circuits.)

The op-amp's two inputs make it ideal for circuits that need a *differential input*. A *differential signal*, also called a *balanced signal*, is composed of two signals of opposite polarity. These two signals are 180° out of phase with each other; when one is at peak positive, the other is at peak negative. When fed signals at both inputs, the op-amp will amplify and pass only the *difference* between those two signals. The two out-of-phase signals will be amplified and combined into one signal at the op-amp's output, but any noise, distortion, or spurious junk common to both polarities will not be amplified by the op-amp. The phenomenon by which a differential amplifier rejects signals common to both channels is called *common-mode rejection*; the op-amp's ability to reject signals common to both channels is called its *common-mode rejection ratio*, or *CMRR*.

The op-amp driven with a differential signal naturally converts a balanced signal to an unbalanced signal. Many audio components with balanced inputs have a differential amplifier at the input, which converts the balanced signal to a single-ended (unbalanced) signal that the component then amplifies.

Digital Electronics

Digital electronics are becoming increasingly important in audio recording and reproduction. Not only CD players and digital processors use digital electronics: more and more control and switching functions in preamps and power amps are being handled by digital circuits.

Simply, digital circuits operate in one of two states: on or off. The “On” state is represented by a voltage (such as +5V), and “Off” is represented by no voltage (also called *ground*). The terms *high* and *low*, or *one* and *zero*, also describe the two conditions in digital circuits. Unlike an analog signal that has a continuously varying amplitude, digital signals are a train of high and low pulses. Digital circuits are also called *binary* circuits to describe the two-state nature of a digital signal.

In digital signals, information is represented as a series of high and low states. Each state carries one piece of information, or *bit*. (Appendix C explains how binary coding can represent numbers, which in turn can represent an audio signal. Binary coding can store virtually any sort of information; this book’s text and graphics were stored as binary code.) The building blocks of digital electronics are called *logic gates*. A logic gate has more than one input, but only one output. Logic gates are made with tiny transistors that are either fully turned on or fully turned off. Simple logic gates can make basic decisions. For example, an *OR gate* will produce a high output if either input #1 *or* input #2 is high. A two-input *AND gate* will produce a high output if input #1 *and* input #2 are high. Digital electronics use many other types of logic gates.

Logic gates can be arrayed to perform nearly any mathematical task. In practice, many hundreds of thousands of gates are needed to make digital circuits. These gates are combined into *integrated circuits*, or ICs. Several hundred thousand gates can be put on a single IC.

Appendix C: Digital Audio Basics

Introduction

Throughout the history of recorded sound, the goal has been to capture an acoustic event by storing a mechanical or magnetic representation of the original acoustic waveform. The modulations in an LP groove, for example, are an analog of the acoustic wave that was originally heard as music. The more alike the groove modulations and the acoustic waveform, the higher the fidelity. Tiny changes in the shape of that squiggly line—inevitably introduced in record cutting, pressing, and playback—produce an almost infinite variability in sound quality.

The advent of digital audio has fundamentally changed the way audio is encoded, stored, and played back. Rather than attempting to preserve the acoustic waveform directly, digital audio converts that waveform into a numerical representation. The numbers representing the waveform are stored and later converted back into an analog signal. Because those discrete numbers can be stored and recovered more precisely than the infinitely variable analog signal, digital audio has the potential to exceed the musical performance of analog media.

I say “potential” because, in many respects, digital audio has yet to sound better than the best analog formats. Virtually anyone who’s listened to a live microphone feed, then the same signal played back both from analog tape and after encoding into digital form, will attest that the analog tape produces a truer representation of the music. Similarly, a comparison of a properly played LP and the CD of the same music reveals musical virtues in the analog representation not heard from the digital medium.

This isn’t to say that analog is without flaw; rather, it suggests that the distortions generated by digital audio are less sonically benign than those from analog media.

Fortunately, this state of affairs is only temporary; digital audio has taken a large step forward with the commercial availability of high-resolution formats

such as SACD and high-res downloads. Only in the last decade did we discover that building digital audio products is a far more exacting process than theory suggests. Although some audiophiles shun digital, it's shortsighted to summarily reject all digital audio based on the crude sound provided by the technologies of digital's first decades. The future promises much more than that.

Binary Number System

All digital audio systems—and digital electronics in general—use the *binary* number system. “Binary” means “two,” reflecting the two possible values in the binary, or base 2, number system: 0 and 1. A binary number of these two values strung together can represent any quantity we want, as can the ten values of the decimal system we’re so familiar with.

When we count in decimal and reach the highest digit (nine), the ones place is reset to zero and the sum is carried to the next column, producing the number 10. This number is a code that represents a 1 in the tens column and a 0 in the 1s column—one ten, no ones. The position of a digit indicates the weight given that digit.

Just as there is no single digit greater than nine in decimal notation, there is no single digit greater than 1 in binary notation. Consequently, when we count in binary, we reset and carry with every count. For example, the number 01 in binary represents the decimal number 1. If we add the binary numbers 01 and 01, we reset the ones column and carry to the next column, producing the binary number 10. This represents a 1 in the twos column and a 0 in the 1s column, or the number 2 (expressed in decimal). Binary addition works like this:

$$\begin{array}{r} 01 \\ +01 \\ \hline 10 \end{array}$$

$$\begin{array}{r} 10 \\ +01 \\ \hline 11 \end{array}$$

$$\begin{array}{r} 11 \\ +01 \\ \hline 100 \end{array}$$

Just as the position of a digit within a decimal number changes its value by a factor of ten (ones, tens, hundreds, thousands, etc.), the position of a binary digit changes its value by a factor of 2. For example, the binary number 0101 indicates a 1 in the 1s column, a 0 in the 2s column, a 1 in the 4s column, and a 0 in the 8s column. The binary number 0101 thus represents the decimal quantity 5. (Fig.C-1 shows binary and decimal equivalences.)

Binary	Decimal
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8
1001	9
1010	10
1011	11
1100	12
1101	13
1110	14
1111	15

Fig. C-1 Binary and decimal equivalences

Any decimal number can be represented by binary notation, provided the binary number is long enough. The highest number that can be represented in binary notation is 2^x , where x is the number of places. For example, a binary number with four places can represent any decimal number between 0 and 15 (2^4). The number 15 in binary is 1111: a one in the 8s column, a one in the 4s column, a one in the 2s column, and a one in the 1s column.

Each of these positions in a binary number is called a *bit*, short for *binary digit*. The greater the number of bits, the higher the number we can represent. The bit in the position of highest value in a binary number (usually at the far left) has the greatest effect on the number's value and is called the *Most Significant Bit*, or *MSB*. The bit in the position of lowest value in the binary number (usually the far right column) is called the *Least Significant Bit*, or *LSB*.

Binary notation makes computers and electronic circuits simpler and more practical. The electronic circuit or device need operate in only two states: one or zero, on or off, yes or no. If a voltage is present, we call that binary 1; if no voltage

is present, that represents binary 0. These conditions relate to a transistor being fully biased on (binary 1) or fully biased off (binary 0). Similarly, a north-south polarization of a magnetic medium could represent binary 1, and a south-north polarization might indicate binary 0.

In digital audio, the analog signal is represented by a stream of these binary numbers. In the next sections, we'll see exactly how music can be represented by binary numbers.

Sampling and Quantization

Sampling and quantization are the cornerstones of digital audio. These two distinct but interrelated processes form the basis for all digital audio systems. To see how sampling and quantization can preserve an analog audio signal as a series of numbers, let's first look at an audio signal's characteristics.

An analog audio signal is a voltage that varies over time. The faster the variation over time, the higher the audio signal's frequency. The greater the amplitude swings, the louder the signal. The audio signal thus has two properties—time and amplitude—that must be encoded to correctly preserve that signal.

An LP record is a good example of how time and amplitude information are preserved. The side-to-side modulations in the record groove encode the amplitude information; the greater the modulation, the higher the signal amplitude. The time information is encoded by the LP's rotation, which must be the same on playback as when the disc was recorded. If we change the speed of the LP, we change the time relationship and thus the audio signal's frequency.

Digital audio must also preserve the time and amplitude information of an audio signal. But instead of encoding and storing those characteristics continuously as in an LP groove, digital audio preserves time and amplitude in discrete units.

A digital audio system encodes the signal's time information by *sampling* the audio signal at discrete time intervals. The amplitude information is encoded by generating a number at each sample point that represents the analog waveform's amplitude at sample time—a process called *quantization*. Sampling preserves the time information, quantization preserves the amplitude information.

Sampling and quantization produce a series of binary numbers—called *words*—that represent the analog waveform. When those binary words are converted back into voltages with the same timing reference as when they were first sampled, the analog audio waveform is roughly reconstructed. The original waveform is recovered by smoothing the resulting staircase shape with a low-pass filter. Sampling and quantization thus convert a continuous analog function (the continuously varying voltage of the analog waveform) into a series of discrete binary numbers. [Fig.C-2](#) shows how a continuous analog waveform is converted into binary numbers, then back into a continuous analog waveform.

Sampling Rate, Nyquist Theorem, and Aliasing

The speed at which the samples are taken is called the *sampling rate*, or *sampling frequency*. The sampling rate determines the highest audio frequency that can be correctly encoded. Specifically, the sampling rate must be at least twice as high as the highest audio signal to be preserved. The compact disc format, for example, samples the analog signal at 44,100 times per second (44.1kHz), which provides an audio bandwidth of about 20kHz. Sampling at 96kHz, provides a bandwidth of nearly 48kHz.

This fundamental relationship between sampling frequency and audio bandwidth was stated in the now-famous Nyquist Theorem. In 1928, telegraph engineer Harry Nyquist postulated that in any sampled system, the sampling rate must be at least twice as high as the highest frequency we want to encode.

If we violate the Nyquist theorem and sample a signal with a frequency higher than half the sampling rate, a form of distortion called *aliasing* occurs. Aliasing creates new frequencies in the sampled signal that weren't present in the original signal. Specifically, if we sample a 33kHz sinewave with a 48kHz sampling frequency, a new signal at 15kHz (48kHz minus 33kHz) is created. Once aliasing occurs, there is no way to remove the spurious alias components from the signal. The term *alias* comes from the fact that the incorrectly sampled signal assumes a completely new identity as a lower-frequency signal. Aliasing is sometimes called "foldback" because the incorrectly sampled high-frequency signal is folded back into the audioband.

The only way to prevent aliasing is to never let the sampler see a signal higher in frequency than half the sampling rate. This is accomplished with an *input low-pass filter* that removes all energy above half the sampling frequency. This input filter, called an *anti-aliasing filter*, is an essential part of all digital audio systems. It allows audioband frequencies to pass, but rapidly attenuates the signal just below half the sampling rate. These filters have been a significant source of sonic degradation in digital audio, particularly in early CD releases, when the filters were more crude than they are today.

Armed with this knowledge that the input signal is band-limited, we can address the common question of newcomers to digital audio: What happens to the information *between* the samples? If there are squiggles in the waveform between samples, won't that information be lost by the sampling process?

The answer is that there is *no* information between samples. If the input signal is properly band-limited, those squiggles will be removed by the anti-aliasing filter. And if the appropriate sampling rate is chosen for the filter's cutoff frequency, the sampler will perfectly encode any waveform input to it. In [Fig.C-2](#), the input filter removed all energy above 20kHz (squiggles in the uppermost sinewave) before the signal was sampled.






Fig.C-2 Graphical representation of the signal flow in a complete PCM digitization system. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)


Let's take the case of a 40kHz sampling frequency and a 20kHz sinewave. Even though the sampler sees only two samples per waveform, that's enough information to perfectly encode that 20kHz sinewave. [Fig.C-3](#) shows the original 20kHz waveform (top) and the resultant waveform after sampling (the squarewave). We know from [Appendix A](#) that a squarewave is a sinewave combined with all its odd-order harmonics; these harmonics are represented as the four high-frequency waveforms below the 20kHz squarewave. When these harmonics are removed by an output low-pass filter, we're left with a 20kHz sinewave—the same signal we started with.




(A) 20 kHz input sine wave




(B) 20 kHz sampled square wave




(C) 180 kHz sine wave (ninth harmonic)



(D) 140 kHz sine wave (seventh harmonic)



(E) 100 kHz sine wave (fifth harmonic)



(F) 60 kHz sine wave (third harmonic)

Fig.C-3 A 20kHz sinewave sampled at 40kHz produces a squarewave. Low-pass filtering removes the harmonics above 20kHz, recovering the 20kHz sinewave. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

We can see from this worst-case example that sampling is a perfectly lossless process: the reconstructed waveform is theoretically identical to the input waveform. This is true provided that the input signal is properly band-limited and an appropriate sampling rate for the filter's characteristics is chosen. Digital audio has its problems (which I'm about to discuss), but sampling theory isn't one of them.

One mechanism, however, by which today's standard-resolution digital audio (CD, for example) degrades musical quality is too low a sampling frequency. Although our upper limit of hearing is 20kHz (if you're over 30, 16kHz is a more accurate figure), research has shown that increasing the audio bandwidth from 20kHz to 40kHz improves musical reproduction. Although we can't hear 40kHz sinewaves, removing energy above 20kHz from the signal reduces the music's sense of openness, transient attack, and natural quality. Moreover, the steep digital filtering required with 44.1kHz sampling degrades sound quality, a phenomenon described in [Chapter 7](#) under "Why 44kHz Sampling and 16-Bit Quantization Isn't Enough."

Quantization

We've seen how sampling preserves the time information in an audio signal; now let's look at how quantization encodes the amplitude information.

Quantization generates a binary number that represents the analog waveform's amplitude at sample time. The binary number is a digital representation of the audio signal's analog voltage when the sample is taken.

The number of bits available to encode the audio waveform's amplitude is called the *quantization word length*. Just as the sampling rate determines the digital audio system's bandwidth, the quantization word length determines the digital audio system's dynamic range, resolution, and distortion level. The most common digital audio systems, including the CD, use 16-bit words. High-resolution digital audio is generally defined as word lengths of 20 bits or more. The longer the word length, the greater the fidelity to the original signal.

Quantization word length defines the number of steps available to encode the audio signal. Specifically, the number of steps is 2^x , where x is the number of bits in the word. For example, 16-bit quantization provides 2^{16} steps (65,536) to quantize the analog signal amplitude. An 18-bit system quadruples this figure to 262,144, and a 20-bit quantization word length provides 1,048,576 quantization steps. The longer the word length, the wider the dynamic range, the lower the distortion and noise, and the higher the resolution.

Unlike sampling, which is a lossless process, quantization produces errors in the encoded signal. Converting an infinitely variable analog function into a discrete number is inherently an approximation. This error occurs because the amplitude represented by a given quantization word is virtually never the signal's exact analog voltage. The difference between the analog signal's actual amplitude and the amplitude represented by the quantization word is called *quantization error*.

Fig.C-4 shows how quantization error occurs. The analog amplitude values don't exactly coincide with the values represented by the quantization words. For example, the first sample (the vertical dotted line to the far left) falls between the quantization levels 100111 and 101000. There is no value "100111.25," so the quantizer simply rounds off to the nearest discrete quantization value (100111), even though that number isn't quite correct. The difference between the voltage represented by 100111 (1.3V) and the audio signal's actual voltage (1.325V) is quantization error. When the audio signal is reconstructed, the slightly wrong

binary number 100111 will generate a slightly wrong analog amplitude. The result is a distortion of the original waveform.




Fig.C-4 Quantization error is limited to one-half LSB. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

The worst case is when the analog amplitude falls exactly between two quantization levels, as occurs at the second sample in Fig.C-4; the difference between the analog amplitude and the quantization word representing that amplitude is the greatest.

Quantization error can be expressed as a percentage of the least significant bit (LSB). At the first sample, the quantization error was one-quarter of an LSB. At the second sample, the error was one-half of an LSB. Note that quantization error can never exceed half the amplitude value of an LSB. Consequently, the smaller the quantization step size, the less the quantization error. Adding one bit to the quantization word length doubles the number of quantization steps and halves the quantization error. Because such a halving represents a difference of 6dB, the signal-to-noise ratio of a digital system improves by 6dB for each additional bit in the quantization word. A digital system with 18-bit quantization will have 12dB less quantization noise than a system with 16-bit quantization.

We can also approximate the signal-to-noise ratio of a digital system by multiplying the number of bits in the quantization word by 6. Sixteen-bit quantization provides a theoretical dynamic range of about 96dB. A 20-bit digital audio system has a dynamic range of about 120dB, or 24dB higher than a 16-bit system.

Quantization error is audible as a rough, granular sound on low-level signals, particularly reverberation decay. Instead of hearing the sound decay into silence, the reverberation decay grows in coarseness and grain as the signal decays. This is because the quantization error becomes a greater percentage of the signal as the signal's amplitude decreases.

This increase in distortion as the signal level drops is unique to digital audio; all analog recording formats exhibit higher distortion at *high* signal levels, such as saturation in analog tape. This situation presents unusual challenges: an increase in distortion as signal level decreases makes that distortion much more audible. Increasing the word length to 20 from 16 bits significantly reduces this problem.

Note that signal-to-noise ratio and distortion figures specified for a digital audio system are with a full-scale signal. Most of the time, a music signal is well below full-scale and therefore closer to the noise floor. The distortion isn't a function of how many bits the system has available, but the number of bits used by the signal at any given moment. This is why distortion and noise are inversely proportional to signal amplitude, and why digital audio has problems with low-level signals.

These factors make setting recording levels in a digital system an art completely different from setting analog recording levels. Ideally, the highest peak over the entire audio program should just reach full-scale digital (using all the bits). If the recording level is set so that the highest peak reaches -6dB , this is equivalent to throwing away one bit of the quantization word, with the attendant 6dB reduction in signal-to-noise ratio. And if the analog signal amplitude is higher than the voltage represented by the highest number, the quantizer simply runs out of bits and repeats the highest number, making the waveform peaks flat. This highly distorted waveform produces an unacceptable "crunching" sound on peaks. If you have a digital audio recording device with meters, you can look at the recording levels on compact discs by taking the digital output from a CD transport and plugging it into the digital recorder. The recorder's meters will show the exact levels recorded on the CD. If the highest peaks never reach full-scale, some resolution has been needlessly lost.

Note that an audio program with a very wide dynamic range will be closer to the quantization noise floor more of the time than a signal with limited dynamic range. The peaks on the wide-dynamic-range-signal will be set just below full-scale, and the much-lower-level signal will consequently be encoded with fewer bits. This is a particular problem in classical music, which has a very wide dynamic range. This is a strong incentive for recording or mastering engineers to compress the music's dynamic range. In addition, producers of pop music want their records to sound loud in relation to other songs over the radio. Severely limiting the music's dynamic range makes the music loud all the time, but at the expense of dynamics, life, and rhythmic power.

Digital signal levels are referenced to a full-scale signal, which occurs when all the bits are used and no higher signal can be coded. This reference level is called 0dBFS, with FS standing for "full-scale." We may refer to a signal as -20dBFS, meaning it is 20dB below full scale.

Dither

The most extreme case of quantization error occurs when the signal amplitude is less than one LSB (Fig.C-5). This low-level signal isn't encoded at all by the quantizer: the quantizer outputs the same code at each sample point and the information is completely lost (C and D). If the signal does traverse quantization steps, the encoded signal is a squarewave, which represents a significant distortion of the original signal. This suggests that any information lower in amplitude than the LSB will be lost.

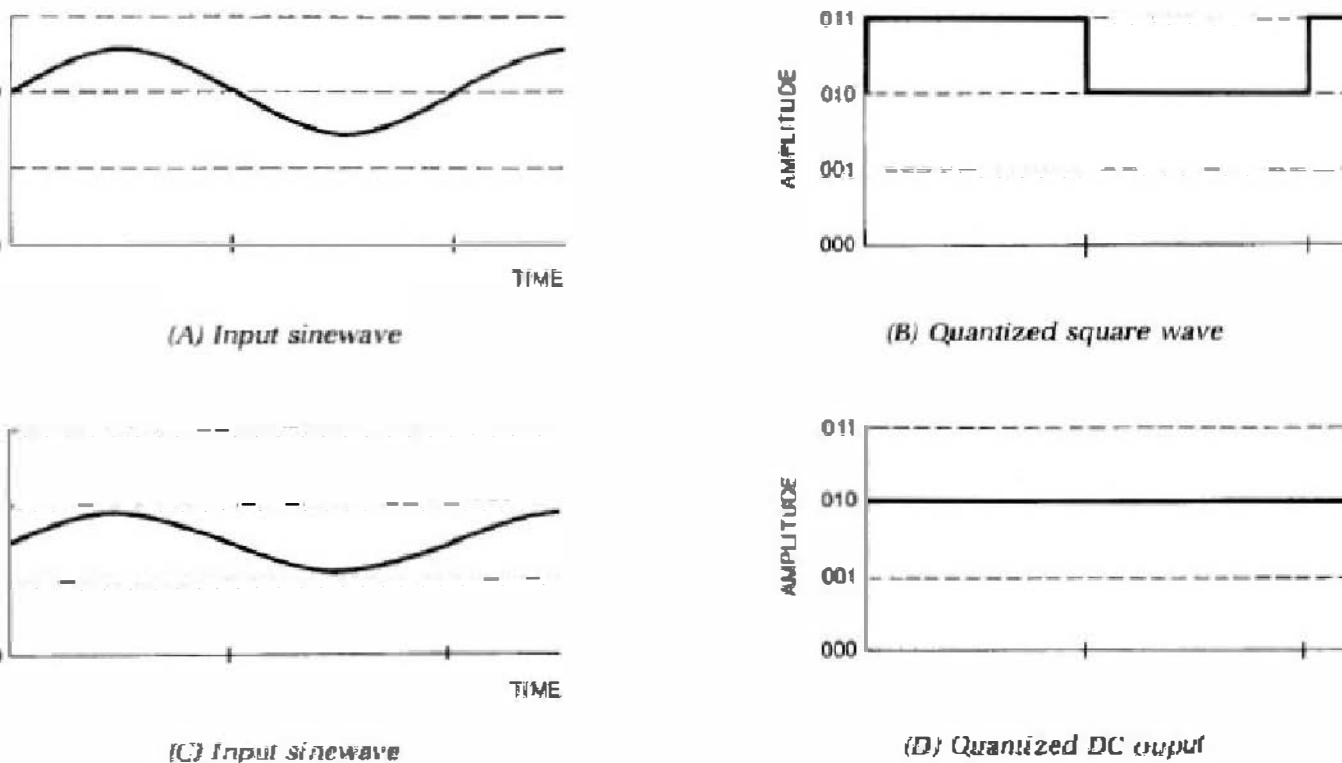



Fig.C-5 Quantization error is large relative to a signal traversing only a few quantization steps. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)


Fortunately, this limit can be overcome by adding a small amount of noise, called *dither*, to the audio signal. Dither allows a quantizer to resolve signals below the LSB, and greatly improves the sound of digital audio.

Fig.C-6 is identical to Fig.C-5, but this time we've added a small amount of white noise to the audio signal. This noise causes the audio signal to traverse


quantization levels, allowing the signal to be encoded. The original sinewave is better preserved in the pulse-width modulation signal of Fig.C-6.




(A) Dither is added to the sinewave input.



(B) Sinewave is preserved in pulse width modulation



(C) Dither is added to the sinewave input.



(D) Sinewave is preserved in pulse width modulation

Fig.C-6 Dither can alleviate the effects of quantization error. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

Dither reduces quantization artifacts, allows the system to resolve information lower in amplitude than one-half the least significant bit, and makes a digital audio system sound more like analog. Among other benefits, dither improves low-level resolution and smooths reverberation decay. Without dither, reverberation decay becomes granular in texture, then seems to drop off into a black hole. It's ironic that a small amount of analog white noise can greatly improve the performance of digital audio. Still, despite the great advantages of dither, we do end up paying a small penalty in decreased signal-to-noise ratio. Dither is an odd concept to someone unfamiliar with digital audio; specifically, that adding noise to the audio signal can improve fidelity.

Finally, it should be pointed out that the preceding explanation of sampling and quantization describes a pulse-code modulation (PCM) digital-audio system,

the standard encoding method in nearly all professional digital audio and in the compact disc and Blu-ray Disc formats. Other forms of encoding are possible such as Pulse-Width Modulation (PWM) used in the Super Audio CD format. This encoding method is explained in [Chapter 7](#).

Digital Audio Storage

Once the audio signal is in the digital domain, we must store it on some medium. Common storage formats include the compact disc, solid-state memory, magneto-optical discs, and computer hard disks.

One might intuitively conceptualize a digital audio recording system as storing the binary information representing the audio signal directly on the medium. For example, a reflective spot on a CD could represent binary one, a non-reflective spot binary zero. In practice, however, digital audio is always encoded with a modulation scheme. Further, a significant amount of additional information is added to the raw audio data.

Every encoding scheme increases the storage capacity of the medium and facilitates data recovery. The Eight-to-Fourteen Modulation (EFM) encoding scheme used in the CD is a good example of encoding. The disc surface is a spiral track of small indentations called *pits* and flat areas called *land*. A transition from pit to land or land to pit represents binary one; all other surfaces (pit bottom or land) represent binary zero. The pattern of pit and land is created by the EFM encoding: 8 audio bits are converted into 14 bits for storage on the CD. These 14-bit words are linked by 3 “merging” bits to produce 17 bits for 8 audio bits. The result is a specific pattern of ones and zeros that follows certain rules. For example, each binary one is separate from other binary ones by a minimum of two zeros and a maximum of ten zeros. Although EFM encoding more than doubles the number of bits to be recorded, it actually increases the storage capacity of the CD compared to unmodulated coding. Note that both left and right audio channels as well as this additional information are encoded into a single serial datastream.

The compact disc is also a good example of how audio data are mixed with nonaudio data and stored on the medium. The CD data-stream is formatted into structures called *frames*. A CD frame (Fig.C-7) contains a synchronization pattern to identify the beginning of a block, bits that contain information about the signal (pre-emphasis, elapsed time, data format, etc.), and error-correction information. Between EFM coding, error correction, subcode, and synchronization, only about a third of the data stored on a CD are actual audio data. All digital audio-storage formats add non-audio data and structure the data into clearly defined blocks.




Fig.C-7 The CD frame format prior to EFM modulation. (Courtesy Ken Pohlmann, *Principles of Digital Audio*)

Error Correction

No storage or transmission medium is perfect. LPs have ticks and pops, and analog magnetic tape has imperfections in the magnetic coating that cause a momentary loss of high frequencies. In analog recording, nothing can be done about these errors; we must live with their sonic imperfections.

Digital audio, however, provides an opportunity to make the audio signal virtually impervious to media imperfections. Lost or damaged data can be reconstructed with error-correction techniques. Error correction doesn't merely approximate the missing data; in most cases it can precisely restore the lost information bit for bit. For example, the compact disc's error-correction system can completely correct up to nearly 4000 consecutive missing bits. These errors could result from poor-quality disc manufacturing, or scratches on the disc from mishandling. Without error correction, the CD simply wouldn't function. The official CD specification—called the "Red Book"—allows a Block Error Rate (BLER) of 220 errors per second.

The two principles on which error correction is based are *redundancy* and *interleaving*. Redundant data is information added to the signal that isn't necessary to decode the signal—unless that signal is corrupted. The crudest form of redundancy is simply to store the data twice. Data corrupted in one spot are unlikely to be corrupted at the second storage location. A more efficient technique transmits a *checksum* with the data that indicates if the data have been reliably retrieved. In practice, error-correction schemes are enormously complex.

Digital audio media are subject to two kinds of errors: *bit errors* and *burst errors*. Bit errors are short in duration and easy to correct. Burst errors, however, affect long, consecutive streams of data. Not only is a large chunk of data missing, but the redundant information accompanying that data is also lost and so the chance of correction is reduced.

A technique called *interleaving* increases the possibility that a burst error can be corrected. Interleaving scatters data to different areas of the media. When those data are deinterleaved (put back in the original order), the long burst error is broken up into many shorter errors that are more easily corrected.

If the error is so large that it cannot be completely corrected, the error is *concealed* by approximating the missing data. Some concealed errors may be audible as a *click* or *pop* sound. This is a rare occurrence; the CD's error-correction system is so good that very few discs—even those abused and scratched—have uncorrectable errors.

Digital-to-Analog Conversion

The digital-to-analog conversion process requires recovering the digital audio data from the medium and converting it to analog form.

The recovered datastream must first be demodulated. In the CD format, this involves converting the 14-bit EFM-coded words back into their original 8-bit form. The data are deinterleaved and error correction performed. Synchronization information—essential to data recovery—is no longer needed and is discarded. The subcode information (track number and timing, along with other housekeeping data) is stripped out of the datastream and handled separately from the audio data. All of these functions are performed by the CD transport, which outputs a formatted datastream to the outboard DAC (or internal digital-to-analog converters in a CD player), which actually converts the signal to analog. (A full description of how a DAC converts digital data to an analog signal is included in [Chapter 7](#).)

Jitter Explained

In analog recording, speed variations in the recording media cause audible distortion. The classic case is a tape machine's "wow and flutter," speed fluctuations that cause the music's pitch to shift up and down slowly ("wow") or quickly ("flutter"). Digital audio doesn't suffer from wow and flutter; digital audio data can be read from the medium at a varying rate, yet be clocked out at a constant speed. But digital audio does suffer from a distortion-producing mechanism conceptually similar to wow and flutter, called *clock jitter*.

Jitter—timing variations in the clocks that serve as the time-base reference in digital audio systems—is a serious and underestimated source of sonic degradation in digital audio. Only recently has jitter begun to get the attention it deserves, both by high-end designers and audio academics. One reason jitter has been overlooked is the exceedingly difficult task of measuring such tiny time variations—on the order of tens of trillionths of a second. Another reason jitter has been ignored is the mistaken belief by some that if the ones and zeros that represent the music are correct, then digital audio must work perfectly. Unfortunately, getting the ones and zeros correct is only part of the equation. Those ones and zeros must be converted back to analog with extraordinarily precise timing to avoid degrading the signal.

As described in [Chapter 7](#), the series of discrete audio samples are converted back into a continuously varying signal with a digital-to-analog converter (DAC) chip. A DAC takes a quantization word and converts it to a voltage—exactly opposite the function of the A/D converter. All that is required for perfect conversion (in the time domain) is that the samples be input to the DAC in the same order they were taken *and with the same timing reference* as when the samples were created. In theory, this sounds easy—just provide a stable clock to the A/D converter and a stable clock of the same frequency to the D/A converter.

It isn't that easy in practice. If the samples aren't converted back to an analog waveform with the identical timing with which they were taken, distortion of the analog waveform will result. These timing errors between samples are caused by variations in the clock signal that controls *when* the DAC converts the digital words to an analog voltage.

Let's take a little closer look at how the DAC decides *when* to convert the digital samples to analog. In [Fig.C-8](#), the binary number at the left is the quantization word that represents the analog waveform's amplitude when it was

first sampled. The bigger the number, the higher the amplitude. (This is only conceptually true—in practice, the data are in 2’s complement form, a code that is easier to handle by circuits.) The squarewave at the top is the “word clock,” the timing signal that tells the DAC *when* to convert the quantization word to an analog voltage. Assuming the original sampling frequency was 44.1kHz, the word clock’s frequency will also be 44.1kHz (or, if the processor uses an oversampling digital filter, some multiple of 44.1kHz). On the word clock’s *leading edge*, the next sample (quantization word) is loaded into the DAC. On the word clock’s *falling edge*, the DAC converts that quantization word to an analog voltage. This process happens 44,100 times per second (without oversampling). If the digital processor has an 8x-oversampling digital filter, the word-clock frequency will be eight times 44,100, or 352.8kHz.




Fig.C-8 The word clock tells the DAC when to convert the quantization words to an analog voltage. (Courtesy *Stereophile*)





Fig.C-9 Word-clock jitter consists of either a random variation or a variation that has a periodic component. (Courtesy *Stereophile*)

It is here at the word clock that timing variations affect the analog output signal. Specifically, clock jitter is any time variation between the clock's trailing edges. **Fig.C-9** shows perfect and jittered clocks (exaggerated for clarity).


Now, look what happens if the samples are reconstructed by a DAC whose word clock is jittered (**Fig.C-10**). The sample amplitudes—the ones and zeros—are correct, *but they're in the wrong places*. The right amplitude at the wrong time is the wrong amplitude. A *time* variation in the word clock produces an *amplitude* variation in the output, causing the waveform to change shape. A change in a waveform's shape is the very definition of distortion. Remember, the word clock tells the DAC *when* to convert the audio sample to an analog voltage; any variations in its accuracy will produce an analog-like variability in the final output signal—the music.

Clock jitter can also raise the noise floor of a digital converter and introduce spurious artifacts. If the jitter has a random distribution (called “white jitter” because of its similarity to white noise), the noise floor will rise. If, however, the word clock is jittered at a specific frequency, artifacts will appear in the analog output as sidebands on either side of the audio signal frequency being converted to analog. It is these specific-frequency artifacts that are the most sonically detrimental; they bear no harmonic relationship to the music, and may be responsible for the hardness and glare often heard from digital audio.

Fig.C-11a is a spectral analysis of a computer-simulated digital processor's output when reproducing a full-scale, 10kHz sinewave with a jitter-free clock. **Fig.C-11b** is the same measurement, but with 2 nano-seconds (two billionths of a second) of jitter with a frequency of 1kHz on the clock. The plot in **Fig.C-11b** reveals the presence of discrete-frequency sidebands on either side of the test signal caused by jitter of a specific frequency. The amplitude of these artifacts is a function of the input signal level and frequency; the higher the signal level and frequency, the higher the sideband amplitude in the analog output signal.





A) Analog signal reconstructed correctly with jitter-free clock.



B) Analog signal reconstructed with jittered clock. (Exaggerated for clarity). Although all the samples have the correct amplitude, because they occur at the wrong time, the reconstructed signal has the wrong shape.

Fig.C-10 An analog waveform is reconstructed correctly with a jitter-free word clock (top); word-clock jitter results in a distortion of the analog waveform's shape. (Courtesy *Stereophile*)



Figs. C-11a and b Computer-simulated spectral analysis of a digital processor's output with a jitter-free clock (left), and with 2ns of jitter (right). (Courtesy *Stereophile*)

How much jitter is audible? In theory, a 16-bit converter must have less than 100 picoseconds of clock jitter if the signal-to-noise ratio isn't to be compromised (there are 1000 picoseconds in a nanosecond). Twenty-bit conversion requires much greater accuracy, on the order of 8ps. One hundred picoseconds is one-tenth of a billionth of a second (10^{-10})—about the same amount of time it takes light to travel *an inch*. Moreover, this maximum allowable figure of 100ps assumes that the jitter is random (white), and does not have a specific frequency that would be less sonically benign. Clearly, extraordinary precision is required for accurate conversion.

Where does clock jitter originate? The primary source is the interface between a CD transport and a DAC. The master clock signal is embedded in the S/PDIF (Sony/Philips Digital Interface Format) signal that connects the two devices (though it's more accurate to say that the audio data are embedded in the clock). The DAC recovers this clock signal at the input receiver chip. The typical method of separating the clock from the data and creating a new clock with a *Phase Locked Loop* (PLL) produces *lots* of jitter. Even if the clock is recovered with low jitter, just about everything inside a digital processor introduces clock jitter: noise from digital circuitry, processing by integrated circuits—even a pcb trace.

It's important to note that the only point where jitter matters is at the DAC's word-clock input. A clock that's recovered perfectly but is degraded just before it gets to the DAC is no better than a high-jitter recovery circuit that's protected from additional jitter on its way to the DAC. Conversely, a highly jittered clock can be cleaned up just before the DAC with no penalty.

Digital Signal Processing (DSP)

DSP is the manipulation of an audio signal in the digital domain. DSP opens vast and exciting new possibilities for improving the quality of reproduced sound—possibilities unimaginable with analog techniques.

DSP is nothing more than the mathematical processing of the digital data representing an audio signal. This processing can take many forms, allowing DSP chips to be configured in a wide range of applications. For example, DSP chips can be programmed to be a high-pass and a low-pass filter, thus replacing the capacitors, resistors, and inductors in a loudspeaker crossover. Some loudspeakers accept a digital input, implement the crossovers with DSP, then convert the signal to analog inside the loudspeaker. The advantages of DSP crossovers are significant: the designer can make the crossover frequency and slope exactly as he or she wants it without regard for the limitations of analog crossover components—capacitors, resistors, and inductors. Moreover, the phase response can be precisely controlled, as can the time offset between drivers. In addition to such control, DSP crossovers get analog components out of the signal path for better sound. Such “digital” loudspeakers were pioneered by Meridian Audio. (DSP-based loudspeakers are covered in [Chapter 6](#).)

Distortions in conventional loudspeakers can be corrected by DSP. If we know that a particular loudspeaker has amplitude errors and departures from linear phase response, those errors can be corrected in the digital domain by DSP.

Perceptual Coding

Storing or transmitting even standard-resolution (16-bit/44.1kHz) digital audio requires a large storage capacity of the recording medium, or fast Internet connection if downloading music. Large storage capacity and wide bandwidth translate to greater expense in delivering digital audio to consumers. This quandary is resolved by reducing the number of bits needed to encode an audio signal. This technique, called *perceptual coding*, attempts to radically reduce the data rate of digital audio signals with minimal effect on sound quality.

Dolby Digital is a good example of why perceptual coding is such a hot field. A 16-bit/44.1kHz digital audio signal (as is found on CD) consumes 705,600 bits per second per channel (16 times 44,100). With Dolby Digital's six discrete audio channels, a conventionally coded film soundtrack would consume a whopping 4,233,600 bits per second (4.23Mbs). This is roughly half the maximum continuous bitrate possible from DVD, leaving very little data for digital video, which has an even more voracious appetite for data capacity. The Dolby Digital perceptual encoder produces a digital audio datastream with a bitrate of just 384,000 or 448,000 bits per second—for all six channels. That's less than one tenth the amount of data without perceptual coding.

Let's look at how this is possible. Perceptual coding exploits idiosyncrasies of human hearing to allocate the available bits to sounds we can hear, and to ignore sounds we're less likely to hear. For example, the psychoacoustic effect of *masking* (described in [Appendix A](#)) prevents us from hearing quiet sounds in the presence of louder sounds. Perceptual coders introduce severe noise and errors into the audio signal because the music is encoded with so few bits. But that noise and distortion are cleverly manipulated so that they are just under the *masking threshold*, or the point below which a sound is inaudible. For example, if most of the musical energy at a given moment is concentrated between 1kHz and 3kHz, the noise will also be concentrated in that frequency band. The noise spectrum is constantly shifting so that the encoding error (noise) "hides" beneath the correctly coded component of the audio signal. Frequency bands that contain more information at any given moment are allocated more bits; those bands with little information are allocated fewer bits. In practice, perceptual coding is vastly more complex.

I was once able to attend a remarkable demonstration of perceptual coding that highlighted how the technology works. I first heard a musical selection that

had a signal-to-noise ratio of 13dB—a ridiculously low figure. For comparison, most electronics have signal-to-noise ratios of 80dB or greater. The sound was so noisy and distorted that the music was barely recognizable. I next heard the same music, again with a signal-to-noise ratio of 13dB, but this time it sounded fine. On an objective basis, the two presentations had exactly the same amount of noise. The difference was that in the second presentation, the noise was constantly shifting in spectral content so that it could “hide” beneath the correctly coded musical signal.

As powerful and useful as this technology can be, it has unfortunately been taken to extremes by some researchers in the field. They pronounced early perceptual coders as “transparent” (the perceptually coded signal was indistinguishable from the original), even though they sounded horrible. Moreover, in their zeal to further reduce the data rates, proponents of perceptual coding have sacrificed musical quality. The researchers speak of “psychoacoustic redundancy” and “informational irrelevance” with little regard for preserving the music with the highest fidelity. With this technological power comes hubris; perceptual coding has been promoted as suitable for master-tape archiving, a truly frightening thought. Our musical heritage should be preserved with the best technology, not the most economical.

Those readers interested in learning more about digital audio are directed to the excellent *Principles of Digital Audio* by Ken Pohlmann.

Glossary

AAC see *Advanced Audio Codec*

A/B comparison A back-and-forth listening comparison between two musical presentations, A and B.

absolute polarity The relationship between the positive and negative-going pressure waves (compressions and rarefactions) of an acoustical event and the positive and negative-going waves on playback. A recording with correct absolute polarity played back through a system with correct absolute polarity will produce a positive pressure wave from the loudspeakers in response to a positive pressure wave at the original acoustic event. Incorrect absolute polarity introduces a 180° phase reversal of both channels. Absolute polarity is audible with some instruments and to some listeners.

AC see *alternating current*

AC line conditioner Device that filters noise from the AC powerline and isolates equipment from voltage spikes and surges. Prevents noise generated by one component in the audio system from getting into other components.

acoustic absorber Any material that absorbs sound; e.g., carpet, drapes, thickly upholstered furniture.

acoustic diffuser Any material that scatters sound, such as a bookcase filled with books of varying depths or commercially made acoustic-treatment devices.

acoustic feedback Sound from a loudspeaker causes a turntable to vibrate; that vibration is converted by the phono cartridge into an electrical signal, which is reproduced by the loudspeaker, which causes the turntable to vibrate even more. This sets up a feedback loop in which the vibration feeds on itself, becoming louder and louder. You've probably heard acoustic feedback at an amplified concert as a howl or screech from the PA system.

acoustic panel absorber Device that absorbs low to mid frequencies by diaphragmatic action. That is, sound striking the panel absorber causes the panel to move, converting acoustic energy into a minute amount of heat in the panel.

acoustics The science of sound behavior. Also refers to a room; e.g., “This room has good acoustics,” meaning the room’s influence on reproduced sound isn’t detrimental to fidelity.

AC synchronous A motor whose rotational speed is determined by the frequency of the AC voltage driving it. Used in most belt-drive turntables.

AC-3 (Audio Codec 3) Another name for *Dolby Digital*, the 5.1-channel discrete digital surround-sound format. (see *Dolby Digital*)

active subwoofer A speaker designed to reproduce only low frequencies and powered by its own integral amplifier.

adaptive USB Type of USB connection in which the receiving device (usually a DAC) locks to the source device’s clock. Compare with *asynchronous USB*, in which the clocks of the receiving and source devices run independently of each other, resulting in lower jitter.

ADC see *analog-to-digital converter*

adjacent-channel selectivity Tuner specification that describes a tuner’s ability to reject the signals of radio stations broadcasting at frequencies adjacent to that of the desired station.

Advanced Audio Codec (AAC) Lossy data-compression file format that reduces the number of bits for transmission or storage. AAC is the default import format for Apple iTunes.

Advanced Audio Distribution Profile (A2DP) One of two lossy data-compression systems used in Bluetooth devices. The other is the superior-sounding aptX. Bluetooth devices support either A2DP or aptX.

AES/EBU interface Professional system for transmitting digital audio. An AES/EBU cable is a balanced line and is terminated with XLR connectors. Also found in some consumer products. Named for the Audio Engineering Society and the European Broadcasting Union.

aggressive Describes a forward sound that seems to thrust the music at the listener.

agile Describes bass reproduction that can accurately reproduce quickly changing pitches and dynamics.

AIFF see *Audio Interchange File Format*

air Describes treble openness, or a sense of space between instruments on the soundstage. Compare with *dull*, *thick*.

air-bearing tonearm A tonearm whose armtube floats on a cushion of air.

air-bearing turntable A turntable whose platter floats on a cushion of air.

AirPlay Trade name for protocols that allow Apple products to communicate and share data wirelessly.

ALAC see *Apple Lossless Audio Codec*

AllMedia Online service that downloads metadata to music servers and streaming services. Formerly called All Media Guide.

alternate-channel selectivity Specification describing a tuner's ability to reject stations two channels away from the desired station.

alternating current (AC) Flow of electrons in a conductor that reverses its polarity many times per second.

ambience Spatial aspects of a film soundtrack that create a sense of size and atmosphere, usually reproduced by the surround speakers. Can also describe the sense of space in two-channel audio-only recordings.

Ambisonics A method of capturing a three-dimensional soundfield and reproducing that soundfield through a multichannel loudspeaker array. Although Ambisonics is compatible with monophonic and stereo reproduction, it is generally thought of as a multichannel music format.

American Wire Gauge (AWG) Standardized method of specifying wire thickness. The lower the AWG number, the thicker the wire.

ampere A unit of electrical current.

amplifier stand An audio accessory that sits on the floor and supports a power amplifier.

analog An analog signal is one in which the varying voltage is an analog of (similar to) the acoustical waveform; i.e., the voltage varies continuously with the original acoustical waveform. Compare with digital signals, in which audio and/or video information is represented by binary ones and zeros.

analog bypass Input on a digital controller that passes along analog input signals to its analog outputs without first converting them from analog to digital and from digital back to analog.

analog-to-digital converter A circuit that converts an analog signal to a digital signal.

analytical Describes the sound of a component that reveals every nuance in the recording, but in an unpleasant way. An analytical component rarely sounds musical.

anechoic Literally, “without echo”; i.e., without resonances.

anechoic chamber A room free of acoustic echoes. An anechoic chamber’s walls are covered in highly absorbent materials that reflect no sound back into the room. Used in loudspeaker testing and measuring.

antiskate adjustment Control on a tonearm that adjusts the amount of force applied to the tonearm to counteract the arm’s tendency to skate (pull inward).

Apple Lossless Audio Codec Usually called Apple Lossless or ALAC. File format that delivers perfect bit-for-bit accuracy to the source data with about a 40% reduction in file size.

aptX One of two lossy data-compression systems used in Bluetooth devices. The other is the inferior-sounding A2DP. Bluetooth devices support either aptX or A2DP.

armtube Long, cylindrical component of a tonearm that runs between the pivot point and the head-shell.

articulate Describes the sound of a component that clearly resolves pitches.

ASIO (Audio Stream Input/Output) Interface protocol for computer audio devices that allows software to directly control the audio device, bypassing the computer’s integral audio-processing functions.

asynchronous sampling-rate converter Sampling-rate converter in which the input and output sampling rates run independent of each other.

asynchronous USB Type of USB connection in which the clocks of the receiving and source devices run independently of each other, resulting in lower jitter. Compare with *adaptive USB*, in which the receiving device locks to the source device’s clock.

atmosphere see *ambience*.

A2DP see *Advanced Audio Distribution Profile*

audioband The range of human hearing; generally considered to be 20Hz–20kHz.

Audio Interchange File Format (AIFF) Digital audio file format in which the data are stored in raw PCM format. Used primarily with Apple computers.

audiophile A person who values sound quality in reproduced music.

Audiophilia nervosa Humorous name for a condition in which the sufferer constantly frets over audio equipment at the expense of his or her enjoyment of the music.

Audio Return Channel Feature of HDMI 1.4a that sends the audio signal back from a television to a controller or A/V receiver. Obviates the need for a separate audio cable when listening through a home-theater system to audio picked up by the television, such as when connected to an over-the-air antenna.

audiostream The flow of digital bits representing an audio signal.

Audio Stream Input/Output see *ASIO*

A/V Abbreviation of audio/video. Identifies a component or system that processes video as well as audio signals.

A/V input An input on an A/V receiver or A/V preamplifier that includes both audio and video jacks.

A/V loop An A/V input and A/V output pair found on all A/V receivers and A/V preamplifiers. Used to connect a component that records as well as plays audio and video signals.

A/V preamplifier A component, aka an A/V controller, that lets you control the playback volume and select which source you want to watch. A/V preamplifiers also perform surround decoding.

A/V preamplifier/tuner An A/V preamplifier that includes, in the same housing, an AM and FM tuner for receiving radio broadcasts.

A/V receiver The central component of a home-theater system; receives signals from source components, is used to select which signal you watch and listen to and to control the playback volume, performs surround decoding, receives radio

broadcasts, and amplifies signals to drive a home-theater loudspeaker system. Also called a surround receiver.

AWG see *American Wire Gauge*

azimuth 1) In magnetic tape recording, the angle of the recording and playback heads referenced to the tape edge (ideally, 90°). 2) In LP playback, the angle of the cantilever in relation to the record surface.

azimuth meter Handheld device that measures a phono cartridge's azimuth indirectly by indicating the left- and right-channel output levels when playing a test LP.

Baerwald alignment Cartridge alignment developed in 1941 by H.G. Baerwald that specifies a certain cartridge offset and overhang for minimal tracking error. One of two optimal cartridge alignments; the other is *Löfgren*.

baffle The front surface of a loudspeaker, on which the drivers are mounted.

balanced-armature driver A tiny transducer found in some in-ear monitors in which a tiny rod (the armature) is balanced like a see-saw within a magnetic field. The audio signal flows through a coil of wire around the armature, magnetizing the armature. One end of the armature is connected to a tiny diaphragm. The magnetic field, varying in response to the audio signal flowing through the coil, pushes and pulls the armature, and with it the diaphragm, to create sound.

balanced connection A method of connecting audio components with three conductors in a single cable. One conductor carries the audio signal, a second conductor carries that audio signal with inverted polarity, and the third conductor is the ground.

balanced headphone drive Method of driving headphones in which each side of the transducer is driven by an amplifier. Compare with unbalanced drive, in which one amplifier is connected to one end of a transducer's voice coil, with the other side of the transducer's voice coil connected to a common ground.

banana jack A small, tubular connector found on loudspeakers and power amplifiers for connecting speaker cables terminated with banana plugs.

banana plug A common speaker-cable termination that fits into a banana jack.

bandwidth The range of frequencies that a device can process or pass. In reference to an electrical or acoustic device, bandwidth is the range of frequencies between the -3dB points.

bass Sounds in the low audio range; generally, frequencies below 500Hz.

bass extension A measure of how deeply an audio system or loudspeaker will reproduce bass; e.g., a small subwoofer may have bass extension to 40Hz, while a large subwoofer may have bass extension to 16Hz.

bass management A combination of controls and circuits in an A/V receiver, A/V preamplifier, or multichannel digital audio player that determines which low frequencies are sent to each loudspeaker and subwoofer.

bass reflex A speaker design with a hole or slot in the cabinet that allows sound inside the cabinet to emerge into the listening room. Bass-reflex speakers have deeper bass extension than speakers with sealed cabinets, but that bass is generally less tightly controlled. Also called “vented” or “ported.” Compare with *infinite baffle*.

BD see *Blu-ray Disc*

bi-amp Using two power amplifiers to drive one loudspeaker. One amplifier typically drives the woofer, the second drives the midrange and tweeter.

binaural recording Recording technique in which microphones are placed inside the “ear canals” of a dummy head. The recording contains spatial cues introduced by the shape of the dummy head and its outer ears, or pinnae. When reproduced through headphones, a binaural recording re-creates the original three-dimensional acoustic space with uncanny realism.

binding post A connection on power amplifiers and loudspeakers for attaching loudspeaker cables.

bipolar speaker A speaker that produces sound equally from the front and the back. Unlike the dipolar speaker, the bipolar’s front and rear soundwaves are in phase with each other.

bipolar transistor The type of transistor most commonly used in audio electronics. The term, which means two poles, comes from fact that current flows through two types of semiconductor material, P and N. Bipolar transistors are either NPN or PNP types, which refers to the polarity of their operating voltages.

bit The smallest unit of binary information. One bit can express one of two states, zero or one. Short for *binary digit*.

bitrate The number of bits per second stored or transmitted by a digital audio or digital video signal. For example, the bit rate of the Compact Disc format is

705,600 bits per second per channel. Higher bit rates generally make possible better audio quality. Also called “datarate.”

bit transparency Ability of a music server to store and output a bitstream that is identical to the bitstream fed to it. A bit-transparent server does not corrupt the audio data.

bi-wire Technique of running two cables to each loudspeaker. One cable is connected to the loudspeaker’s woofer input terminals, one to the tweeter’s input terminals. Bi-wiring is possible only with loudspeakers with two pairs of input terminals.

bi-wire jumpers Short pieces of wire or metal that connect the two pairs of binding posts on a bi-wirable loudspeaker when the loudspeaker is connected to the amplifier by a single run of wire. One jumper connects the two positive (red) terminals to each other, and a second jumper connects the two negative (black) terminals to each other.

bleached Describes the sound of a component that emphasizes the upper harmonics of an instrument’s sound while de-emphasizing the lower harmonics and fundamental. A bleached sound is thin, bright, and lacks warmth.

bloom A sense of air around instrumental images.

Bluetooth Wireless communication system that allows audio to be streamed between devices.

Blu-ray Disc (BD) 120mm optical disc format developed as the high-definition successor to the DVD. Can store high-resolution multichannel audio in addition to high-definition video. An audio-only version can deliver up to eight channels of PCM audio at 192kHz/24-bit.

boom Excessive bass, particularly over a wide band of frequencies.

Bostik Blu-Tack Putty-like material, available in hardware stores, that can be applied between a stand-mounted loudspeaker and its stand, to damp vibrations.

break-in Initial period of use of a new audio component, during which time the component’s sound might improve.

bridge rectifier An arrangement of four diodes that converts the alternating current (AC) from the powerline into direct current (DC).

bridging Method of connecting a stereo amplifier to a loudspeaker that converts a 2-channel amplifier into a 1 -channel power amplifier. One amplifier channel

amplifies the positive half of the waveform, the other channel amplifies the negative half. The loudspeaker is connected as the “bridge” between the two amplifier channels.

brightness Excessive treble that makes the sound more shrill.

brittle A midrange or treble character that makes instrumental timbres sound harsh. Compare with *liquid*.

buffer Electronic circuit that isolates one component or circuit stage from another. A preamplifier is a buffer because it acts as an intermediary between source components and a power amplifier; the preamplifier’s buffering function relieves the source components of the burden of driving a power amplifier. A digital buffer temporarily stores data in solid-state memory; digital buffers act as jitter-reduction devices because their output is controlled by a precision clock.

build quality An audio or video component’s quality of electronic parts, chassis, and construction techniques.

burn To record digital files on an optical disc.

burn in Process of improving an audio component’s sound quality by leaving it powered on, or by running audio signals through it.

burn-in signal A specialized signal that burns in a component more rapidly than does a music signal.

bypass test Listening test in which the component under evaluation is alternately removed from and reinserted in the signal path, in order to assess its sonic characteristics.

byte A group of eight bits.

cable elevator Device that raises a cable off the floor.

cable enhancer Device designed to break in a cable by sending a high-level audio signal through the cable.

calibration The act of fine-tuning an audio or video component for more accurate performance. In an audio system, calibration includes setting the individual channel levels.

cantilever Thin tube protruding from a phono cartridge to which the stylus is attached.

capacitive reactance Property of capacitance that blocks low frequencies but passes along high frequencies. Capacitive reactance makes a capacitor behave as a frequency-dependent resistor. Because of capacitive reactance, a capacitor connected to a tweeter allows treble to pass but blocks bass.

capacitor Electronic component that stores a charge of electrons. Reservoir capacitors are used in power amplifiers to store energy; filter capacitors filter traces of AC from DC power supplies; coupling capacitors connect one amplifier stage to another by blocking DC and allowing the AC audio signal to pass.

capture ratio Tuner specification: the difference in dB between the strengths of two stations required before a tuner can lock to the stronger station and reject the weaker one. The lower the capture ratio, the better the tuner.

cartridge-alignment protractor Device for precisely aligning a phono cartridge within a tonearm's headshell.

cartridge demagnetizer Device that removes stray magnetic fields from metal parts inside a phono cartridge.

CD (Compact Disc) A polycarbonate disc 120mm in diameter that stores up to 80 minutes of music.

CD-4 (Compatible Discrete 4) A Quadraphonic format invented by JVC in which the two rear channels are encoded along with the front channels; a high-frequency (30–50kHz) difference signal allows the rear-channel signals to be demodulated.

CD Recordable (CD-R) Format for recording digital audio on a compact disc. CD-R is a “write once” format: once recorded on, the disc cannot be erased.

CD-Rewritable (CD-RW) Recordable CD format that can be rerecorded multiple times.

center channel In a multichannel audio system, the channel that carries information that will be reproduced by a speaker placed in the center of the viewing angle, between the left and right speakers. The center channel carries nearly all of a film's dialogue.

center-channel mode A setting on A/V receivers and A/V preamplifiers that configures the receiver or preamplifier for the type of center-channel speaker in the system.

center-channel speaker The speaker in a home-theater system, placed atop, beneath, or behind the visual image; reproduces center-channel information, such as dialogue and other sounds associated with onscreen action.

channel balance The relative levels or volumes of the channels in an audio system or individual component.

channel separation A measure of how well sounds in one channel are isolated from sounds in the other channels.

chassis The framework on which an audio component is built. The case typically fits over the chassis.

Chesky 6.0 An alternative multichannel audio technique proposed by Chesky Records using four front and two rear loudspeakers.

chesty A loudspeaker coloration that emphasizes the deep chest sounds in spoken or singing voice.

chuffing Sound created by the port of a bass-reflex loudspeaker when reproducing low bass at high levels; caused by the rapid movement through the port of large amounts of air.

circumaural (literally, “around the ear”) A type of headphone in which the ear cushion surrounds the entire outer ear. Compare with supra-aural headphones, which rest directly on the ears.

Class-A Mode of amplifier operation in which a transistor or tube amplifies the entire audio signal.

Class-A/B Mode of amplifier operation in which the output stage operates in Class-A at low output power, then switches to Class-B at higher output power.

Class-B Mode of amplifier operation in which one tube or transistor amplifies the positive half of an audio signal, and a second tube or transistor amplifies the negative half.

Class D Mode of amplifier operation in which the output transistors rapidly turn fully on and fully off. The analog waveform is encoded in the widths of the resulting pulses. See also *switching amplifier*.

client In networked audio, a device that receives content from a server.

clipping An amplifier that is asked to deliver more power than it is capable of will flatten the tops and bottoms of the audio waveform, making the peaks appear to

be clipped off when the waveform is viewed on an oscilloscope. Clipping introduces a large amount of distortion, audible as a crunching sound on musical peaks.

clock A signal present in all digital products that acts as a timekeeper to synchronize the processing of digital data. Can also refer to an outboard device that provides a precision clock to a DAC or disc transport.

coaxial cable A cable in which an inner conductor is surrounded by a braided conductor that acts as a shield.

coaxial digital output A jack found on CD transports and some CD players that provides a digital audio signal on an RCA jack for connection to another component through a coaxial digital interconnect.

coaxial driver A loudspeaker driver in which one drive unit (typically a tweeter) is mounted inside a second drive unit (typically a midrange). Aka a concentric driver.

codec (compression/decompression) Device or software for converting a digitized audio signal into a format for transmission or storage (compression), and then converting the coded signal back into its original form (decompression).

coherence The impression that the music is an integrated whole, rather than made up of separate parts.

coincident driver Loudspeaker driver in which a smaller driver is mounted within a larger driver. Also called a *coaxial driver*.

coincident microphones A stereo or multichannel recording technique in which microphones are placed very close to each other. Examples include Blumlein (also called crossed figure-8), XY (crossed cardioids), and MS (middle-side: a forward-oriented omnidirectional or cardioid microphone and a figure-8 microphone).

coloration A change in sound introduced by a component in an audio system. A loudspeaker that produces colored sound doesn't accurately reproduce the signal fed to it; e.g., a speaker with coloration may have too much bass and not enough treble, or vice versa.

comb filtering A phenomenon that introduces a series of deep, narrow bands of attenuation in an audio signal, usually caused by combining a direct signal with a slightly delayed version of that signal, such as the reflection by a listening room's sidewall of soundwaves produced by a loudspeaker.

common-mode rejection When a balanced signal is input to a differential amplifier, only the difference between the two phases of the balanced signal is amplified. Any noise common to both phases (common-mode noise) is rejected by the differential amplifier.

compliance In phono cartridges, a number expressing the cantilever's stiffness. Specifically, compliance is the length of cantilever movement when a force of 10^{-6} dynes is applied, expressed in millionths of a centimeter. High compliance is any value above 20. Low compliance is any value below 10.

compliant feet Small devices made from a rubber-like material, usually Sorbothane, that, when placed under an audio component, isolate the component from vibration.

compression A component of a soundwave in which the density of the air is greater than the density of the surrounding air. A woofer cone moving outward creates a compression. One compression and one rarefaction make one complete cycle of a soundwave. Compare with *rarefaction*.

cone The diaphragm of a loudspeaker (made of paper, plastic, metal, or some exotic material) which moves back and forth to set air in motion and thus create sound.

congested A thickening of the sound that makes instrumental images less separate and distinct.

constrained-layer damping Construction method in which a layer(s) of soft material(s) is placed between layers of harder materials to dissipate vibrational energy.

Consumer Electronics Show (CES) Annual trade show for the consumer-electronics industry, held in Las Vegas in early January.

contact cleaner Fluid applied to audio plugs, jacks, cable terminations, and AC cords to remove oxidation and increase conductivity.

contact enhancer Fluid applied to audio plugs, jacks, cable terminations, and AC cords to increase conductivity.

continuousness Quality of reproduced sound in which the music seems a unified whole, rather than disjointed.

control components For example, an audio system's preamplifier, because it is used to adjust the volume and select which source to listen to. Compare with *source components* and *playback components*.

controller Another term for an A/V preamplifier.

critical listening The art of evaluating audio equipment quality by careful analytical listening for specific sonic flaws.

crossfeed circuit Headphone technology that electronically simulates the interaural crosstalk that occurs when listening to loudspeakers, reducing the "inside-the-head" imaging effect of listening through headphones. A crossfeed circuit adds a small amount of left-channel information to the right channel, and vice versa.

crossover A circuit that divides the audioband into two or more frequency ranges. Crossovers are found in virtually all dynamic loudspeakers, and in A/V receivers and A/V preamplifiers.

crossover distortion see *zero-crossing distortion*

crossover frequency The frequency at which a crossover divides the audioband. A subwoofer with a crossover frequency of 80Hz filters all information above 80Hz out of the signal driving the subwoofer, and all information below 80Hz out of the signal driving the main speakers.

crossover slope Describes the steepness of a crossover filter, expressed as "xdB/octave." For example, a subwoofer with a crossover frequency of 80Hz and a slope of 6dB/octave would begin attenuating signals with a frequency higher than 80Hz, with that attenuation increasing with frequency so that the attenuation is 6dB at 160Hz (an octave above 80Hz). A slope of 12dB/octave produces a more rapid attenuation above the crossover frequency, attenuating the signal by 12dB at 160Hz. The most common crossover slopes are 12dB/octave, 18dB/octave, and 24dB/octave. Crossover slopes are also referred to as "first-order" (6dB/octave), "second-order" (12dB/octave), "third-order" (18dB/octave), and "fourth-order" (24dB/octave). The "steeper" slopes (such as 24dB/octave) split the audioband more narrowly and produce less overlap between the two frequency ranges.

crosstalk see *channel separation*

current The flow of electrons in a conductor. For example, a power amplifier “pushes” electrical current through speaker cables and a speaker’s voice coils to make the coils move back and forth.

current-to-voltage converter Circuit following the digital-to-analog converter chip that converts the DAC’s current pulses to a voltage.

cutoff frequency see *crossover frequency*

DAC see *digital-to-analog converter*

damp To decrease the amplitude and duration of vibration.

damping device Accessory component that reduces vibration.

damping factor A number that expresses a power amplifier’s ability to control woofer motion. Related to the amplifier’s output impedance.

DAS see *direct attached storage*

data compression see *perceptual coding*

datarate see *bitrate*

dB see *decibel*

dBFS (decibels referenced to full scale) In digital audio, the maximum recorded level possible in digital. All digital levels are expressed in dBFS; a signal level 20dB below full scale is -20dBFS.

dB SPL (decibels referenced to sound-pressure level) A measurement, in decibels, of loudness (sound-pressure level) referenced to the threshold of hearing (0dB SPL).

dBW (decibel watt) A measurement of amplifier power in decibels (dB) referenced to one watt (W) of output power.

DC see *direct current*

decibel (dB) The standard unit for expressing relative power or amplitude levels of voltage, electrical power, acoustical power, or sound pressure.

dedicated AC line An AC circuit that powers only your audio equipment.

de-emphasis In the Compact Disc format, a treble boost can be applied during recording (*pre-emphasis*), and a treble cut during playback (*de-emphasis*), to increase the signal/noise ratio.

depth The impression of instruments, voices, or sounds being layered or placed behind one another in three dimensions, as in “soundstage depth.”

detail Low-level components of recorded music, such as the fine inner structure of an instrument’s timbre.

dialogue intelligibility The ability to clearly hear and understand, without strain, the dialogue in a film’s soundtrack. Dialogue intelligibility is affected by the quality of components in a home-theater system, room acoustics, and the positions of the speakers.

diaphragm The surface of a loudspeaker driver that moves to set the air in motion and create sound.

dielectric Insulating material inside a capacitor or around a cable.

differential amplifier Electronic circuit that amplifies only the difference between the two phases of a balanced signal.

diffraction The bending of soundwaves as they pass around an object. Also: the re-radiation of sound caused by discontinuities in surfaces near the radiating device, such as the bolts that secure drivers to a speaker cabinet.

diffusion Scattering of sound. Diffusion reduces the listener’s ability to determine the location of the source of a sound, which increases the sense of space produced by surround loudspeakers.

digital Calculation or representation by discrete units. For example, digital audio and digital video can be represented by a series of binary ones and zeros.

digital amplifier Amplifier that accepts a *pulse-code-modulated* (PCM) audio signal and converts it to a *pulse-width-modulated* (PWM) signal that turns the amplifier’s output transistors on and off. Not to be confused with a *Class-D* amplifier, which has an analog signal path and a switching output stage.

Digital eXtreme Definition (DXD) PCM format with a sampling frequency of 352.8kHz and a word length of 32 bits.

Digital Living Network Alliance (DLNA) Communication protocols in consumer-electronics products that allow DLNA-compliant products to recognize each other and share data without user intervention.

digital loudspeaker Loudspeaker incorporating a digital crossover and power amplifiers. A digital speaker takes in a digital bitstream, uses digital signal processing to divide the audioband into several frequency bands, converts the

resulting signals to analog, and amplifies each separately. Each driver is powered by its own amplifier.

digital preamplifier Audio component that performs source selection, volume control, and other functions in the digital domain.

digital signal processing (DSP) Manipulation of audio or video signals by performing mathematical functions on the digitally encoded signal.

Digital Theater Systems (DTS) A discrete digital surround-sound format used in movie theaters and some home-theater systems. A better-sounding alternative to Dolby Digital, the DTS format encompasses film soundtracks and music releases. Aka DTS Digital Surround.

digital-to-analog converter (DAC) 1) A chip that converts digital audio signals to analog audio signals. 2) A standalone component in an audio system that converts digital audio signals to analog signals. Aka a digital processor.

Digital Transmission Content Protection (DTCP) Copy-protection system that allows content owners to control how their content may be copied; aka "5C," for the five companies that developed the standard: Hitachi, Intel, Matsushita, Sony, Toshiba.

digital volume control Digital circuit that adjusts the signal level by performing mathematical computation on the ones and zeros representing the audio signal.

diode Electronic component that allows current to flow in only one direction.

dip A reduction in energy over a band of frequencies. Compare with *peak*.

DIP (dual inline package) switch Small grouping of switches on some phono preamplifiers that provide adjustment for amplifier gain and cartridge loading.

dipolar speaker A loudspeaker that produces sound from the rear as well as from the front, with the front and rear sounds out of phase with each other.

direct attached storage (DAS) Type of storage device (usually a hard-disk drive) in which the drive is addressable only by the host computer. Compare with *network attached storage (NAS)*, in which the storage device can be addressed by other devices on the network.

direct-coupled Electronic circuit in which no coupling capacitors are in the signal path.

direct current (DC) Flow of electrons that remains steady rather than fluctuating. Compare with *alternating current (AC)*.

Direct Stream Digital (DSD) Method of digitally encoding music with a very fast sampling rate but with only 1-bit quantization. Developed by Sony and Philips for the Super Audio CD (SACD).

Direct Stream Digital over PCM (DSD-over-PCM, or DoP) Method of transmitting DSD audio data over USB by disguising it as PCM data.

direct-to-disc Recording format in which the LP master lacquer is cut in real time as the music is performed, with no intermediate storage medium.

discrete Separate. A discrete circuit uses separate transistors instead of an integrated circuit. A discrete digital surround-sound format contains 5.1 channels of audio information that are completely separate from each other; this differs from a matrixed surround format such as Dolby Surround, in which the channels are mixed together for transmission or storage.

dispersion The directional pattern over which a loudspeaker distributes its sound.

dither A small amount of noise added to an audio signal that improves the low-level resolution of digital audio.

DLNA see *Digital Living Network Alliance*

do-it-yourself (DIY) music server A personal computer adapted to function as a music server. Compare with *turnkey music server*.

Dolby Digital A 5.1-channel discrete digital surround-sound format used in movie theaters and consumer formats. The surround format required by the DVD format and for HDTV broadcasts.

Dolby Digital EX Surround format that matrix-encodes a third surround channel in the existing left and right surround channels of a Dolby Digital signal. This third surround channel (called “surround back”) drives a speaker(s) placed directly behind the listening position, for greater spatial realism. Called THX EX in THX-certified products.

Dolby Digital Plus Surround format that provides a higher bitrate than Dolby Digital along with more efficient coding, both of which contribute to higher sound quality. Dolby Digital Plus is an extension of Dolby Digital in that it appends additional data to a Dolby Digital signal. This feature makes Dolby

Digital Plus backward-compatible with Dolby Digital decoders. Dolby Digital Plus was developed for the Blu-ray Disc.

Dolby Pro Logic A type of Dolby Surround decoder with better sound than standard Dolby Surround decoding. Specifically, Pro Logic decoding provides greater channel separation and a center-speaker output. A Dolby Pro Logic decoder accepts a 2-channel, Dolby Surround-encoded audio signal, then decodes that signal into separate signals for the left, center, right, and surround channels.

Dolby Pro Logic II Pro Logic surround decoder with better performance and more features than older Pro Logic circuits.

Dolby Pro Logic IIx Pro Logic surround decoder that generates the surround-back channel from 2- or 5.1-channel sources.

Dolby Pro Logic IIz Surround decoding format that provides outputs that drive additional height loudspeakers located above the left and right front speakers.

Dolby Surround Encoding format that combines four channels (left, center, right, surround) into two channels for transmission or storage. On playback, a Dolby Pro-Logic decoder separates the two channels into four channels.

Dolby TrueHD Surround format developed for Blu-ray Disc. Unlike Dolby Digital, which is a lossy compression format, TrueHD is lossless in that it delivers high-resolution multichannel audio with perfect bit-for-bit fidelity to the source data.

DoP see *Direct Stream Digital over PCM*

double-DSD Direct Stream Digital format that runs at twice the sampling frequency of standard DSD. Also called DSD128 because double-DSD runs at 5.6448 megabits per second, which is 128 times the Compact Disc format's sampling rate of 44,100 kilobits per second.

downconvert Translate a digital audio signal from a higher to a lower sampling frequency; e.g., a digital master tape recorded at 96kHz is downconverted to 44.1kHz for storage on a CD.

download (verb) In computer-based music: to access music from the Internet and store the file for later playback. A download (noun) is a music file accessed from the Internet.

downmix converter A circuit found in DVD players that converts a 5.1-channel discrete Dolby Digital soundtrack into a 2-channel Dolby Surround-encoded

signal. A DVD player's downmix converter lets you hear surround sound from DVD if you don't have a Dolby Digital decoder. Also refers to the circuit in all DVD-Audio players that creates, on the fly, a 2-channel mix from the disc's multichannel signal.

downsample To reduce the sampling frequency of a digital audio signal. Aka "downconversion."

drive To send an audio or video signal from one device to another; e.g., a power amplifier drives a loudspeaker.

driver An individual transducer that converts an electrical signal into mechanical motion, creating sound. A moving-coil driver includes a diaphragm, voice coil, spider, basket, and surround. Most loudspeakers comprise several drivers, each of which is designed to reproduce a specific portion of the audioband.

driver mounting bolts Bolts or screws that attach the drivers to the loudspeaker's baffle (front panel).

dry Sound in which the amount of ambience or reverberation is reduced. Can describe the live sound of a venue, the sound of a recording, or the sound of a system and/or listening room.

DSD see *Direct Stream Digital*

DSD-over-PCM see *Direct Stream Digital over PCM*

DSP see *digital signal processing*

DSP room correction Technique of removing room-induced frequency-response peaks and dips with digital signal processing.

DTCP see *Digital Transmission Content Protection*

DTS see *Digital Theater Systems*

DTS ES Discrete Surround-sound format that delivers discrete 6.1-channel sound. The additional channel drives a loudspeaker or loudspeakers placed directly behind the listening position. Compare with 5.1-channel sound.

DTS ES Matrix Surround-sound format that matrix-encodes a third surround channel into the existing left and right surround channels of a DTS signal.

DTS HD Surround format developed for the Blu-ray Disc that provides a higher bitrate and more efficient coding than DTS, both of which improve sound quality.

DTS HD Master Audio Surround-sound format developed for the Blu-ray Disc that provides high-resolution multichannel audio with perfect bit-for-bit fidelity to the source data. “Master Audio” signifies lossless compression.

DTS Neo:6 Decoding circuit developed by Digital Theater Systems for creating 6-channel audio from 2-channel sources. Comes in two varieties, one optimized for reproducing music (DTS Neo:6 Music) and one optimized for reproducing film soundtracks (DTS Neo:6 Cinema).

dull Sound lacking in treble energy.

DVD 120mm disc that can contain digital video (DVD-Video), high-resolution digital audio (DVDAudio), or computer data (DVD-ROM).

DVD-Audio Subset of the DVD format standard that provides for high-resolution, multichannel digital audio on DVD.

DVD-Video Format created for storing movies on a 120mm disc.

DXD see *Digital eXtreme Definition*

dynamic compression Phenomenon in loudspeaker drivers in which an increase in drive signal results in little increase in acoustic output. Occurs when the driver’s voice coil heats up, which increases the coil’s resistance.

dynamic range In audio, the difference in volume between the maximum and minimum signal levels in a given performance or recording. In video, the difference in light level between black and white (also called contrast).

dynamic-range compressor A circuit found in some Dolby Digital-equipped receivers and preamplifiers that reduces the difference between the maximum and minimum volume levels. A dynamic-range compressor can reduce the volume of peaks, increase the volume of low-level sounds, or both. Useful for late-night listening, when you don’t want explosions to disturb other family members but still want to hear low-level sounds clearly.

dyne A unit of force. The threshold of hearing (0dB SPL) represents a pressure on your eardrum of 0.0002 dynes/cm². A force of 10⁻⁶ dynes applied to a phono cartridge’s cantilever will push that cantilever approximately 15 millionths of a centimeter.

earbud A small transducer worn in the ear. An earbud is pressed against the opening of the ear canal; compare with an in-ear monitor, which is inserted *in* the ear canal.

earphones A pair of small transducers that fit inside each ear canal and reproduce sound.

eartip The pliable cushion surrounding the transducer in earphones and in-ear monitors that creates a seal in the ear canal between the transducer in the ear canal and the outside air.

edge A harshness in transient musical events, such as the initial attacks of percussion instruments.

effective tonearm mass The total mass of a tonearm's moving parts, and where along the tonearm that mass is distributed. Mass near the pivot point only slightly increases a tonearm's effective mass, but the same amount of mass near the tonearm's cartridge end greatly increases the effective mass. A tonearm's effective mass must be matched to the phono cartridge's compliance for optimal sound quality.

electromagnetic interference (EMI) Disturbance caused by the radiation of electromagnetic waves.

electrostatic driver Transducer used in loudspeakers or headphones in which a thin diaphragm is moved back and forth by forces created by the attraction and repulsion between two electrically charged surfaces. Compare with a dynamic loudspeaker, in which a diaphragm is moved back and forth by forces created by magnetic repulsion and attraction.

EMI see *electromagnetic interference*

envelopment The impression of being surrounded by sound, such as with a surround-sound system.

equalization In tape or LP record playback, a treble cut to counteract a treble boost applied during recording. Also describes the modification of a system's tonal balance by employing an equalizer.

equalizer A circuit that changes the tonal balance of an audio program. Bass and treble controls are types of equalizers.

etch An unpleasant emphasis of transient musical information accompanied by excessive brightness.

Ethernet Interface for connecting computers and Ethernet-equipped audio components to a local area network (LAN).

extension How high or low in frequency an audio component can reproduce sound.

Fast Fourier Transform (FFT) Mathematical technique that converts information from the time domain to the frequency domain. Widely used in loudspeaker measurements; a sound impulse created by a loudspeaker is recorded and then, via FFT, converted into a frequency-response plot.

feedback In amplifier circuits, returning part of the output signal to the amplifier's input. Feedback reduces distortion and makes the circuit more stable. See also *acoustic feedback*.

FFT see *Fast Fourier Transform*

file-based music Accessing music that has been digitally encoded into files, such as in computer audio.

filter Electronic circuit that selectively removes or reduces the amplitude of certain frequencies.

filter capacitor Large capacitor used in power supplies to filter traces of 60Hz AC from the DC supply.

FireWire A high-speed interface that can carry high-resolution digital audio. Officially called IEEE1394, but also known as i.LINK (Sony's trade name). An open standard developed by Apple.

5C see *Digital Transmission Content Protection*

5.1 channels The standard number of channels for encoding film soundtracks. The five channels are the left, center, right, surround left, and surround right. The ".1" channel is a 100Hz-bandwidth channel reserved for high-impact bass effects reproduced by a subwoofer or subwoofers.

FLAC (Free Lossless Digital Audio Codec) Digital audio file format that provides perfect bit-for-bit accuracy to the source through lossless compression. A FLAC file is approximately 60% the size of a WAV file of the same data.

flat A speaker that accurately reproduces the signal fed to it is called "flat" because that is the shape of its frequency-response curve. Flat also describes a soundstage lacking in depth.

floorstanding speaker A speaker that sits on the floor rather than on a stand.

flutter echo Back-and-forth acoustic reflections in a room between untreated parallel surfaces (think of a pair of facing mirrors, each reflecting an image into the other). Flutter echo can be heard as a “pinging” sound after a handclap.

forward A sonic presentation in which sounds seem to be projected toward the listener.

frequency Number of repetitions of a cycle. Measured in hertz (Hz), or cycles, per second. An audio signal with a frequency of 1000Hz (1kHz) undergoes 1000 cycles of a sinewave per second.

frequency response A graphical representation of an audio device’s relative amplitude as a function of frequency.

full-range speaker A speaker that reproduces low bass as well as midrange and treble frequencies.

gain Number expressing the amount of amplification provided by an amplifier. An amplifier that can convert a signal with an amplitude of 0.1V to 1V is said to have a gain of 10.

geometry In cables, the physical arrangement of the conductors and dielectric.

gigabit switch A small box in a networked audio system that connects multiple devices to the network so that they can share the same Internet connection. Also called an “Ethernet switch.”

gimbal bearing Tonearm bearing in which the arm is held within a pivoted support that allows the rotation of the tonearm around a single axis. Compare with *unipivot*.

Gracenote CDDB Database of album titles, song titles, and artist names that can be embedded in digital audio servers, computer-based audio systems, and portable music players, or accessed via the Internet.

grain Roughness in instrumental or vocal timbres.

ground A reference point from which all voltages are measured; typically, the potential of the earth.

ground isolator Small device, usually a transformer, that disconnects a component’s ground, to prevent a ground loop.

ground loop In AC power, unwanted current flow through the ground caused by multiple ground paths of varying resistances.

half-speed mastering Technique of cutting an LP record in which the master tape and the lathe's turntable run at half the normal speed, resulting in correct pitch and tempo when played at normal speed. Half-speed mastering gives the cutting head more time to trace the signal, resulting in better sound.

harmonic distortion The production of spurious frequencies at multiples of the original frequency. A circuit amplifying a 1kHz sinewave will create frequencies at 2kHz (second harmonic), 3kHz (third harmonic), and so forth.

HDCD see *High Definition Compatible Digital*

Head-related transfer function (HRTF) Modification of sound by the human torso, head, and outer ears.

headshell Component of a tonearm, located at the end of the armtube, that holds the phono cartridge.

heatsink Large metal device that draws heat away from the interior of an electronic device to dissipate it in air. The fins protruding from the sides of power amplifiers are heatsinks.

Helmholtz resonator Acoustic device with a small opening that absorbs sound by causing the air in the device to resonate at a particular frequency.

hertz (Hz) The unit of frequency; the number of cycles per second; kilohertz (kHz) is thousands of cycles per second.

hierarchical encoding In Ambisonics, a multichannel signal that is scalable to any number of playback channels, including mono and stereo, ensuring compatibility with a wide range of playback systems.

high-blend circuit FM tuner feature that automatically switches the signal to monaural when the broadcast signal's strength falls below a certain level.

high-cut filter Filter that removes treble energy.

High Definition Compatible Digital (HDCD) Process for improving the sound quality of 44.1kHz/16-bit digital audio on Compact Disc. An HDCD-encoded disc is playable in any CD player, but sounds best when played by a CD player or digital processor equipped with an HDCD decoder.

high-density layer The information layer in a hybrid Super Audio CD that contains high-resolution digital audio.

high-pass filter A circuit that allows high frequencies to pass but blocks low frequencies. Also called a “low-cut filter.” High-pass filters are often found in A/V receivers and A/V preamplifiers to keep bass out of the front speakers when a subwoofer is used. Many powered subwoofers also include high-pass filters.

high-resolution digital audio Generally regarded as digital audio with a sampling rate of 88.1kHz or higher and a word length longer than 16 bits.

home theater The combination of high-quality sound and video in the home.

Home THX A set of patents, technologies, and playback standards for reproducing film soundtracks in the home. THX doesn't compete with but builds on surround formats such as Dolby Digital and DTS.

HRTF see *head-related transfer function*

hybrid An audio component that combines more than one technology; e.g., an amplifier that contains tubes and transistors, a loudspeaker that contains dynamic and ribbon drivers, etc.

hybrid disc SACD-based disc playable by CD and SACD players. Has two information layers: one layer contains CD-quality audio, the second contains high-resolution audio.

Hz see *hertz*

IC (integrated circuit) A semiconductor that incorporates many parts, such as transistors and resistors, in one monolithic package, aka a chip.

i.LINK Sony trade name for the IEEE1394 (FireWire) interface.

image specificity The precise positioning and definition of instrumental or vocal images on a recorded or reproduced soundstage.

imaging The impression of hearing, in reproduced music, instruments and voices as objects in space.

IMD see *intermodulation distortion*

immediate A quality of sound that is vivid, lively, and forward. Compare with *laid-back*.

impedance Resistance to the flow of AC electrical current. An impedance is a combination of resistance, inductive reactance, and capacitive reactance, measured in ohms.

inductive reactance The property of an inductor in which the inductor's opposition to current flow increases as frequency increases.

inductor Electronic component that increases its opposition to current flow as frequency increases. Many loudspeaker crossovers use a large inductor on the woofer to filter out high frequencies.

in-ear monitor Transducer that fits inside the ear canal, and whose shape has been custom molded for one specific person's ear canal.

infinite baffle A sealed loudspeaker cabinet. The cabinet wraps around the drive units, mimicking a baffle of infinite size by completely isolating from each other the front- and rear-firing waves produced by the diaphragm. Compare with a *reflex*, or ported, loudspeaker.

infrared (IR) The frequency of light emitted by wireless remote controls.

in phase Stereo speaker connection in which each speaker's positive (red) terminal is connected to the amplifier's positive terminal, and each speaker's negative (black) terminal is connected to the amplifier's negative terminal. When fed a monophonic signal, both speakers' cones move forward and backward in synchrony. Contrast with *out-of-phase* connection in which one speaker's cone moves forward while the second speaker's cone moves backward. Speakers connected out of phase have diffuse imaging and reduced bass output.

input impedance The resistance to current flow presented by a circuit or component to the circuit or component driving it. Impedance is a combination of resistance, capacitive reactance, and inductive reactance.

input overload A condition in which a component (typically, a preamplifier) produces severe distortion because the signal driving it is too high in level. For example, connecting a high-output moving-magnet cartridge to a preamplifier's moving-coil input could overload the preamplifier's input and produce audible distortion.

input receiver Chip or circuit in a digital processor that takes in the S/PDIF datastream, generates a new digital clock based on that datastream, strips out the audio data, and sends the audio data to the digital filter. Converts an S/PDIF bitstream to audio data, control data, and a clock.

integer upsampling Increasing the sampling frequency of a digital audio signal by a whole number; e.g., 44.1kHz to 88.2kHz (x2) or 176.4kHz (x4). Compare with non-integer upsampling; e.g., 44.1kHz to 96kHz (x2.1768707).

integrated amplifier Audio product combining a preamplifier and power amplifier on a single chassis in a single case or housing.

interaural crosstalk Phenomenon that occurs when listening to loudspeakers in which some left-channel information is heard by the right ear and some right-channel information by the left ear. The lack of interaural crosstalk when listening through headphones leads to “inside-the-head” imaging.

interconnect A cable that carries line-level audio signals (audio interconnect) or S/PDIF-encoded digital audio (digital interconnect).

intermodulation distortion (IMD) Distortion introduced by amplifier circuits that comprises the sum and difference of the input signal. For example, an amplifier driven with a mix of 1kHz and 5kHz frequencies generates intermodulation-distortion products at 6kHz (the sum of 1kHz and 5kHz) and at 4kHz (the difference between 1kHz and 5kHz). These IMD products then intermodulate with each other to create a nearly infinite series of distortion products.

Internet Radio Distribution of radio-station broadcast programs via the Internet. Also called “webcasting.”

interpolation Filling in missing audio and/or video information with a “best-guess” estimate based on the information that came before and after the missing section.

iOS Apple operating system for mobile devices such as the iPhone and iPad.

isolation transformer A transformer whose only function is to transfer signals or power from one point to another without a physical connection. It does this through electromagnetic induction, in which two coils of wire are positioned next to each other. Current flow through one coil induces a current in the second coil.

I²S Pronounced *I squared s*. Digital audio interface in which the clock and audio signals are carried on separate conductors.

JFET (junction field effect transistor) Type of transistor often used in preamplifier input stages that differs in internal structure and operation from the more common bipolar transistor. JFETs are lower in noise and are driven by voltage rather than current, more closely mimicking the operation of vacuum tubes.

JIS (Japanese Industrial Standard) plug Small five-pin connector found on many tonearms.

jitter Timing variations in the clock that synchronizes events in a digital audio system. The clock could be in an analog-to-digital converter that controls when each audio sample is taken. Of more interest to audiophiles is clock jitter in digital audio reproduction; the clock controls the timing of the reconstruction of digital audio samples into an analog signal. Jitter degrades sound quality.

jitter-reduction device A component inserted between a digital source, such as a CD transport, and a digital-to-analog converter. The jitter-reduction device accepts a jittered S/PDIF signal and outputs a low-jitter S/PDIF signal.

kbs (kilobits per second) Thousands of bits per second.

kHz see *kilohertz*

kilobit One thousand bits, abbreviated “kb.”

kilobyte One thousand bytes, abbreviated “kB.”

kilohertz One thousand hertz, abbreviated “kHz.”

laid-back Describes a sound in which the music has a sense of ease, with some distance between the listener and the soundstage—as if the listener were sitting in a distant row in a concert hall. Compare with *forward* and *immediate*.

lean Describes a sound that lacks midbass. Synonyms: *thin*, *lightweight*, *underdamped* (to describe loudspeakers). Compare with *weighty*, *full*, *heavy*.

least significant bit (LSB) In binary notation, the bit that has the lowest value. In 16-bit digital audio, the 16th bit has the least effect on the numerical value of the 16-bit word. An example in decimal notation: in the decimal number 15,389, the digit in the ones column (9) has the least effect on the value of the number. Compare with *most significant bit (MSB)*.

LED (light emitting diode) Solid-state device that emits light when current is passed through it. LEDs are used as status indicators on the front panels of many audio components.

level matching Ensuring that components under evaluation are reproducing a recording at precisely the same volume level.

LFE see *low frequency effects*

light-emitting diode see *LED*

lightweight Sound that lacks bass energy.

linear An output that varies in precisely direct proportion to the input. Can also describe a traditional (non-switching) power supply.

linearity Specification of how much a digital-to-analog converter's output level differs from its input level. For example, a DAC driven with a digital signal at -90dBFS that reproduces that level as -88dB has a linearity error of 2dB .

line level An audio signal with an amplitude of approximately 1V to 2V . Audio components interface at line level through interconnects. Compare with *speaker level*, the much more powerful signals that drive loudspeakers.

line-source loudspeaker Loudspeaker with a tall, narrow dispersion pattern. A tall ribbon driver is naturally a line source, as is a vertical array of point-source drivers. Compare with a point-source loudspeaker, which has a shorter, broader dispersion pattern.

liquid Describes a sound that lacks shrillness. Usually applied to the midrange, *liquidity* implies correct reproduction of musical timbres.

litz cable Cable in which many small strands are grouped, with each strand separately insulated to isolate it from the adjacent strands.

loading Method of enclosing a driver (usually the woofer) in a loudspeaker cabinet. Different loading methods include the *infinite baffle*, *bass reflex*, and *transmission line*.

localization The ability to detect the directionality of sounds. Although grammatically incorrect, *localization* is the accepted term in the audio literature.

Löfgren alignment Cartridge alignment developed in 1938 by E. Löfgren that specifies a certain cartridge offset and overhang for minimum tracking error. One of two optimal alignments; the other is *Baerwald*.

lossless compression In digital storage, reducing the number of bits without losing any information. *FLAC* and *ALAC* are examples of lossless compression. Compare with *lossy compression*.

lossy compression Method of reducing the number of bits representing an audio signal by discarding information. *MP3* and *AAC* are examples of lossy compression. Contrast with *lossless compression*.

loudspeaker Device that converts an electrical signal into sound. The loudspeaker is the last link in the playback chain.

low-cut filter A circuit that removes bass frequencies from an audio signal. Also called a *high-pass filter*.

low-frequency cutoff The point at which a loudspeaker's output drops in the bass by 3dB.

low frequency effects (LFE) A separate channel in the Dolby Digital and DTS formats reserved for low-bass effects, such as explosions. The LFE channel is the “.1” channel in the 5.1-channel or 7.1-channel Dolby or DTS format. The LFE channel is usually reproduced by a subwoofer.

low-pass filter A circuit that removes midrange and treble frequencies from an audio signal. Also called a *high-cut filter*.

LP long-playing record Double-sided, 12” vinyl disc that stores approximately 25 minutes of audio per side.

LSB see *least significant bit*

lush see *liquid*

luxury products Audio components built with lavish cosmetic design and expensive housings. Compare with *value products*, which are designed to produce the best sound for the least money.

mass loading Technique of increasing the mass of loudspeaker stands or equipment racks by filling them with sand, lead shot, or other dense materials.

matrix A method of encoding four audio channels into two for transmission or storage.

Mbs Mega (million) bits per second.

MDF see *medium-density fiberboard*

medium-density fiberboard (MDF) Composite wood material from which most loudspeaker cabinets are made.

megabit One million bits, abbreviated “Mb.”

megabyte One million bytes, abbreviated “MB.”

Meridian High Resolution (MHR) Proprietary high-resolution, multichannel, encrypted digital interface used between products made by Meridian Audio.

Meridian Lossless Packing (MLP) Data-compression system used in DVD-Audio that reduces the bit rate with no loss in sound quality. The decompressed bitstream is bit-for-bit identical to the original bitstream.

metadata Non-audio data attached to a music file that contains artist name, track time, musical genre, and other information.

metallic Unpleasantly harsh sound reminiscent of metal being struck.

MHR see *Meridian High Resolution*

microphony Phenomenon in which electrical components convert mechanical vibrations into electrical signals.

midrange Audio frequencies in the middle of the audioband, such as the human voice. Generally, the range of frequencies from about 300Hz to about 2kHz. Also: a driver in a loudspeaker that reproduces the range of frequencies in the middle of the audioband.

millisecond One one-thousandth of a second: 0.001s, 1/1000 second.

minimonitor A small, stand-mounted loudspeaker.

MLP see *Meridian Lossless Packing*

mold-release compound Material used in LP manufacturing to allow the just-pressed record to be easily removed from the stamper. Mold-release compound remains on the surfaces of uncleaned LPs, and can degrade sound quality.

monoblock A power amplifier with only one channel.

MOSFET (metal-oxide semiconductor field effect transistor) A transistor that is turned on (biased) by voltage rather than by current. Compare with *bipolar transistor*.

most significant bit (MSB) In binary notation, the bit that has the highest value. In 16-bit digital audio, the first bit has the greatest effect on the numerical value of the 16-bit word. An example in decimal notation: in the decimal number 15,389, the digit in the ten-thousands column (1) has the most effect on the value of the number. Compare with *least significant bit (LSB)*.

moving-coil cartridge Transducer that converts the motion of the stylus in the record groove into an electrical signal. Tiny coils attached to the cantilever are moved back and forth in a fixed magnetic field, inducing current flow through the coils.

moving-magnet cartridge Transducer that converts the motion of the stylus in the record groove into an electrical signal. Tiny magnets attached to the cantilever are moved back and forth between fixed coils of wire, inducing current flow through the coils.

MP3 (MPEG-1 Audio Layer 3) Low-bit-rate coding-and-decoding scheme used in portable audio players and some Internet downloads.

MSB see *most significant bit*

multibit DAC Digital-to-analog converter employing a “ladder” of resistors, each resistor corresponding to one bit of the digital audio word. The value of each successive resistor is twice that of the resistor on the rung below it, which allows half as much current flow as the resistor on the rung above. This resistance-doubling progression follows the binary progression 1, 2, 4, 8, 16, 32, 64, etc. Compare with a 1-bit DAC, in which the digital audio signal is represented by pulses just one bit high.

multichannel power amplifier A power amplifier with more than two channels; usually, five or seven.

multichannel preamplifier A preamplifier with more than two channels; usually, six or eight.

multichannel sound Sound reproduction via more than two channels feeding more than two loudspeakers.

multiroom A feature on some A/V products that lets you listen to two different sources in two different rooms.

music server Device that stores, organizes, and provides access to digitally encoded music files.

NAS see *network attached storage*

nasal Describes a coloration of the midrange, particularly in loudspeakers, that makes singers sound congested.

nearfield listening Listening to loudspeakers positioned only a short distance away; e.g., listening to desktop speakers.

network attached storage (NAS) Type of storage device (usually, a hard-disk drive) that is connected to a local area network (LAN) and can be addressed by other devices on the network. Compare with *direct attached storage (DAS)*, in which the drive can be accessed by only the host computer.

networked audio An audio system in which the components are connected to each other via a local area network (LAN).

new old stock (NOS) Vacuum tubes manufactured long ago—perhaps decades ago—that have never been used.

noise shaping Digital audio technique in which quantization noise is shifted to the top octave of the audioband or above the audioband, thus decreasing the audible noise. Used in 1-bit DACs and Direct Stream Digital.

NOS see *new old stock*

Nyquist theorem Expresses relationship between a digital audio system's sampling frequency and the highest audio signal that can be preserved. Specifically, the sampling frequency must be at least twice as high as the highest audio frequency to be recorded. Named for telegraph engineer Harry Nyquist, who published the theorem in 1928.

octave The interval between two frequencies with a ratio of 2:1. The bottommost octave of the audioband is 20–40Hz; the topmost octave is 10–20kHz.

off-axis response A loudspeaker's frequency response, measured at the loudspeaker's sides. Compare with the *on-axis response*: a loudspeaker's response directly in front of the speaker's baffle.

offset A bend in a pivoted tonearm that gives the arm an S or J shape and orients the cartridge at 25° relative to the arm. The position of the cartridge relative to the record groove.

Ogg Vorbis Lossy data-compression scheme and file format used by some consumer-electronics products.

ohm Ω The unit of resistance to electrical current flow.

Ohm's law Expresses the relationship between voltage, current, and resistance in electrical circuits. If you know two of the values (e.g., current and resistance), you can calculate the third value using Ohm's law. Named for its discoverer, Georg Ohm.

omnipolar loudspeaker Trade name of Mirage Loudspeakers' proprietary loudspeaker technology for radiating sound in all directions.

on-axis response A loudspeaker's frequency response measured directly in front of the speaker's baffle.

one-note bass Reproduced bass that seems to have just one pitch. Caused by a loudspeaker's excessive output over a narrow band of low frequencies. A "boom truck" is one-note bass taken to an extreme.

op-amp (operational amplifier) An integrated circuit that combines several transistors and supporting circuitry into a ready-made amplifier. Compare with a *discrete* amplifier comprising individual transistors.

opaque Sound that is cloudy, veiled, or thick. Compare with *transparent*.

out of phase Stereo speaker connection in which one speaker is connected to the amplifier with the positive (red) and negative (black) wires reversed. When fed a monophonic signal, one speaker's cone moves forward while the second speaker's cone moves backward. Speakers connected out of phase have diffuse imaging and reduced bass. Contrast with *in-phase* connection, in which both speakers' cones move forward and backward in synchrony.

output impedance Technically, the change in a component's output voltage in response to a change in load impedance. Practically, a component with a high output impedance can deliver less current to a load than a component with a low output impedance. Low output impedance is preferable.

output stage The last amplifier circuit in an audio component. A DAC's output stage is an amplifier that drives the preamplifier. A power amplifier's output stage delivers current to the loudspeakers.

output transformer Transformer in tubed amplifiers that couples the output stage to the loudspeaker. Output transformers are required in tubed amplifiers to change the amplifier's high output impedance to a lower value that can better drive loudspeakers. The output transformer also blocks DC from appearing at the amplifier's output terminals.

overall perspective Describes the degree of forwardness of a recording and/or audio system. An *upfront* perspective is similar to sitting in row C of a concert hall; a *laid-back* perspective is similar to sitting in row W.

overdamped Loudspeaker sound characterized by articulate but lean bass. Compare with *underdamped*.

overhang In loudspeakers, the degree of movement of the woofer after the drive signal has ceased. In turntables, the distance between the tonearm's pivot point and the stylus.

oversampling digital filter Digital audio circuit that removes all energy above half the sampling frequency. An 8x-oversampling filter generates seven new samples for each incoming sample, for an eightfold increase in the sampling frequency. Oversampling shifts the spurious images of the audio signal generated by the sampling process to a higher frequency, where they are more easily filtered out.

pace The sense of rhythmic drive or propulsion in the reproduction of recorded music.

palpable Impression that the reproduced sound of an instrument or voice is tangible; literally, that it can be touched.

pan The continuous movement of sounds and images from location to location. Originally a camera term, from *panorama*.

passband Range of frequencies allowed through by a filter.

passive level control Device that adjusts the volume of an audio system without using amplifiers or other active electronic components, such as tubes or transistors.

passive radiator Loudspeaker diaphragm that isn't connected electrically, but is moved solely by air pressure inside the cabinet created by the motion of the woofer(s); aka an auxiliary bass radiator (ABR).

passive subwoofer A speaker for reproducing bass frequencies that must be powered by a separate power amplifier. Compare with *active subwoofer* or *powered subwoofer*, which contain an integral amplifier.

PCM see *pulse code modulation*

peak A short-term, high-level audio signal. Also: an excess of energy over a narrow frequency band (compare with dip).

peaky Sound with excessive energy over only a narrow frequency band.

pentode Vacuum tube with five elements: the cathode, plate, screen grid, control grid, and suppressor grid. More efficient than the simpler triode vacuum tube.

perceptual coding A method of reducing the number of bits needed to encode an audio or video signal by omitting information unlikely to be heard or seen. MP3 is an example of perceptual coding.

period The time it takes to complete one cycle of a sine wave.

phantom center-channel mode A setting on A/V receivers or A/V controllers invoked when no center-channel speaker is used.

phantom image The creation of an apparent sound source between two loudspeakers.

phase In a periodic wave, the fraction of a cycle that has elapsed, measured in degrees, with 360° representing one cycle. Describes the time relationship between two signals. For example, two sinewaves, the second of which is delayed half a cycle relative to the first, are 180° out of phase.

phase adjustment A control provided on some subwoofers that lets you slightly delay the subwoofer's sound output so that its output is coincident in time with the output from the main speakers.

phase angle A measurement of how inductive or capacitive a loudspeaker's impedance is.

phase-locked loop (PLL) A electronic circuit that compares the phase of an incoming signal with that of a reference frequency.

phone jack A $\frac{1}{4}$ " jack that accepts a $\frac{1}{4}$ " phone plug.

phone plug A $\frac{1}{4}$ "-diameter plug found on full-size headphones. Phone plugs can be mono (one conductor, one ground) or stereo (two conductors, one ground).

pinch effect In LP playback, a pushing up of the stylus caused by a narrowing of the groove during large-amplitude high-frequency signals, caused by the different shapes of the cutting and playback styli.

pit and land The tiny indentations (pits) structures on the surface (land) of an optical disc (CD, DVD, SACD, BD, in which the digital data are encoded.

pitch definition The ability to distinguish pitch in reproduced music, particularly in the bass. Some products, especially loudspeakers, obscure the different pitches of individual notes.

pivoted tonearm A tonearm in which the cartridge and armtube traverse the record in an arc around a fixed pivot point. Compare with *tangential-tracking tonearm*.

planar loudspeaker Loudspeaker in which the driver(s) are mounted in an open frame rather than in a cabinet.

planar-magnetic driver Loudspeaker or headphone driver in which conductors carrying the audio signal are bonded to a diaphragm. A subset of the ribbon driver, the planar-magnetic driver is also called a quasi-ribbon.

platter The “table” part of a turntable; it turns to spin the LP.

platter bearing The bearing on which the platter turns.

platter mat Thin, flat disc that rests atop the platter and under the LP, to better couple the LP to the platter.

playback components Broadly speaking, the power amplifier and loudspeakers. Compare with *source components* (tuner, CD player, LP turntable, Blu-ray Disc player, etc.) and *control components* (preamplifier).

plinth The flat deck of a turntable beneath the platter, aka the “base.”

PLL see *phase-locked loop*

pneumatic platform Device that isolates components from vibration with compressed air.

point-source loudspeaker Loudspeaker that emits sound from a single point in space. Compare with *line-source loudspeaker*.

port Opening in a loudspeaker cabinet that channels bass frequencies from inside the enclosure to outside the enclosure; aka a “vent.”

ported loudspeaker see *reflex-loaded loudspeaker*

port noise Noise generated by the rapid movement of air through the port of a reflex-loaded loudspeaker; aka “chuffing.”

power amplifier An audio component that boosts a line-level signal to a signal powerful enough to drive loudspeakers.

power bandwidth The range of frequencies over which an amplifier can deliver its rated power output.

powered subwoofer A subwoofer with an integral amplifier that drives the subwoofer’s cone.

power handling How much amplifier power, measured in watts, a speaker can accept before it is damaged.

power output A power amplifier’s ability, measured in watts, to deliver electrical voltage and current to a speaker.

power supply Circuitry, found in every audio component, that converts 60Hz alternating current (AC) from the wall outlet to direct current (DC) to drive the audio circuitry.

power transformer Device in a power supply that reduces the incoming voltage from 120V to a lower voltage.

preamplifier Component that receives signals from source components, selects a source for listening, controls the volume, and drives the power amplifier; literally, “before the amplifier.”

preamplifier-processor Audio/video component that combines a preamplifier and surround-sound decoder on a single chassis. Can be thought of as an A/V receiver without power amplification or a tuner.

pre-emphasis In the Compact Disc format, a treble boost can be applied during recording (*pre-emphasis*), and a treble cut during playback (*de-emphasis*), to increase the signal/noise ratio.

pre-preamplifier Component that boosts a tiny signal from a phono cartridge to a higher level for input to a phono preamplifier. Typically used to match a moving-coil cartridge’s output to a preamplifier designed for a moving-magnet input.

presence A quality of reproduced sound that gives the listener the sense that an instrument or voice is actually in the listening room.

presence region Band of frequencies in the midrange (1–3kHz) that contributes to *presence*.

psychoacoustics The study of human hearing.

pulse-code modulation (PCM) Digitally representing an audio signal with a series of samples, each of which typically contains 16–24 bits.

pulse width modulation (PWM) A digital representation of an audio signal in which the signal’s amplitude information is contained in the widths of a single-bit pulse.

punchy A quality of sound with dynamic impact, particularly in the bass.

push-pull amplifier Amplifier in which pairs of transistors or vacuum tubes are arranged so that one transistor or tube amplifies the positive half of the audio waveform, and the other transistor or tube amplifies the negative half. Compare with *single-ended amplifier*.

PWM see *pulse width modulation*

Q or Quality factor Numerical value describing woofer resonance in an enclosure. Technically, a woofer's Q is its peak *resonant frequency* divided by the peak's bandwidth. A high-Q speaker is said to be underdamped; its bass is less tight and defined than that of a low-Q, or overdamped, speaker.

Quadraphonic Four audio channels. Two primary Quadraphonic formats were introduced in the early 1970s: the matrix SQ system and the discrete CD-4.

Quad-rate Direct Stream Digital (DSD) Format with a bitrate four times that of standard DSD. Also called DSD256 because the sample rate is 11.2MHz, or 256 times the Compact Disc's sample rate of 44,100Hz.

quantization Assigning a discrete numerical value to an analog function. In digital audio, the analog waveform's amplitude is converted to a number (quantized) each time a sample is taken.

quantization error Difference between the actual analog value and the number representing that analog value. Quantization error occurs when the analog value falls between two quantization steps; the quantizer assigns the closest number. Quantization error introduces noise and distortion in digital audio, often heard as a roughness at low signal levels, particularly during reverberation decays.

quantization word length The number of bits created by the A/D converter at each sample point. Compact Disc records quantization words that are 16 bits long.

quarter-wavelength bass trap Acoustical absorber in which an absorbent material is hung at a distance from a reflective surface. The peak absorption of a given frequency occurs when the distance between the absorbent material and the reflective surface is one-fourth the wavelength of sound, and at odd multiples of a quarter wavelength.

quasi-ribbon driver see *planar-magnetic driver*

quieting sensitivity Tuner specification that states the voltage across the antenna required to produce an audio signal with a signal/noise ratio of 50dB. Compare with the less stringent *usable sensitivity*.

radiation pattern The width and height through which a loudspeaker disperses sound.

Radio Data System (RDS) Standardized protocol for broadcasting text or other data along with an FM radio broadcast. The receiving tuner must be equipped with RDS decoding to display the text or data.

radio frequency (RF) noise Most often refers to spurious noise with energy in the radio-frequency band of several hundred kilohertz to tens of megahertz.

RAID see *redundant array of independent disks*

rail The power-supply voltage that supplies a component's audio circuits.

rail fuse Fuse sometimes found on the back of a power amplifier, wired in series with the rail. The rail fuse will blow when excessive current is drawn by the amplifier's output stage.

rarefaction Component of a soundwave in which the density of air is lower than that of the surrounding air. A woofer cone moving inward creates a rarefaction. One rarefaction and one compression make one complete cycle of a soundwave.

RCA jack A connector found on audio and video products. Analog and digital signals transmitted via RCA jacks include line-level audio and S/PDIF digital.

RDS see *Radio Data System*

rear wave The sound radiated from the back of a loudspeaker, away from the listener. The sound radiated by the back of a driver into the speaker enclosure.

record brush Device used to remove dust and debris from the surface of a record.

record clamp Device that couples a record to the platter, through weight or via a threaded screw.

record flattener Device that gently heats a record, then presses it flat.

record sleeve The paper or plastic liner inside a cardboard LP jacket.

rectifier Device that converts AC current into DC current. Found in all power supplies.

rectify To convert AC current into DC current.

Red Book The original published specifications for the Compact Disc format; named for that document's red cover. A "Red Book" disc is a conventional CD.

Red Book layer The information layer of a Super Audio CD that contains conventional 44.1kHz/16-bit digital audio.

redundant array of independent disks (RAID) Data-storage technique in which multiple hard disks are operated as a unit, with one or more disks storing a duplicate of the data. In a typical two-disk array, the second disk stores a duplicate copy of the first disk's data that can be used as a backup if the first disk fails.

re-equalization A Home THX technology that reduces the amount of treble on playback so that you hear a more natural sound when a film soundtrack is played in the home.

reflex-loaded loudspeaker see *bass reflex*

renderer In networked audio, a device that receives a digital audiostream from a server and converts it to an analog signal.

resolution The ability of an audio component to reveal low-level musical information.

resonance Vibration of an object or air that is disproportionate in amplitude to the stimulus. A bell ringing at a certain frequency is an example of resonance; the bell's pitch is its resonant frequency.

resonant frequency Frequency at which an object, or the air in a room, naturally vibrates when stimulated.

reverberation Dense acoustical reflections in an acoustic space that become lower in amplitude and repeat more rapidly. The sound in a room after the source of the sound has stopped producing sound.

reverberation time (RT) The time it takes sound in a room to decrease in amplitude by 60dB, expressed as RT_{60} . A small living room may have an RT_{60} of 0.3 second; a large cathedral might have an RT_{60} of 5 seconds.

RF filter Device that blocks *radio-frequency noise* from passing through a cable.

RIAA Recording Industry Association of America.

RIAA accuracy Flatness of the RIAA phono-equalization circuit in a phono preamplifier.

RIAA equalization A treble boost and bass cut applied to the audio signal when a record is cut; a reciprocal treble cut and bass boost on playback restores flat response. RIAA equalization makes it possible to increase the playing time of an LP side (bass takes up the most space in the record groove) and decrease noise (the treble cut on playback also reduces record-surface noise).

ribbon loudspeaker Loudspeaker in which the diaphragm is electrically conductive and carries the audio signal. Usually made from a long, thin strip of aluminum.

rip (verb) To transfer data from a storage medium (e.g., a CD or LP) to a writable medium such as a hard-disk drive. (noun) A file created by copying it from another storage medium.

ripping NAS A *network attached storage* (NAS) device that contains an optical disc drive that can read data from optical discs and store it on the NAS.

ripping software Third-party software that reads data from a CD for storage on a writable medium.

rolled off Sound with reduced energy at the frequency extremes (bass and/or treble). A loudspeaker with rolled-off treble sounds dull.

rolloff Reduction in energy at the frequency extremes, or the effect of a filter; e.g., “the crossover produces a 12dB/octave rolloff above 2kHz.”

room correction see *DSP room correction*

room gain Increase in apparent bass output when a loudspeaker plays in a room compared with that speaker’s bass output in an anechoic chamber. The room’s walls increase the amount of bass perceived; the closer to the walls the speaker is placed, the greater the room gain.

room resonance modes Excessive acoustical energy at certain frequencies when the air in a room is excited by the sound of a loudspeaker.

rumble Unwanted very-low-frequency noise in LP playback produced by the turntable’s motor, bearing, or other subsystems.

SACD see *Super Audio CD*

sampling The process of converting an analog audio signal into digital form by taking periodic “snapshots” of the audio signal at some regular interval. Each snapshot (sample) is assigned a number that represents the analog signal’s amplitude at the moment the sample was taken.

sampling frequency The rate at which samples are taken when recording digitally. Expressed in samples per second or, more commonly, in hertz. A digital recorder working at 144.1kHz/16-bit, the standard resolution of the Compact Disc, samples the music 44,100 times per second.

sampling-rate converter Circuit that changes the sampling frequency of a digital audio signal.

satellite speaker A small speaker with limited bass output, designed to be used with a subwoofer.

screechy Unpleasant character of the upper midrange or treble that makes instruments or voices sound like noisy car brakes. Most often applied to the sound of violins.

secondary windings Coils of wire in a transformer that aren't physically connected to the AC line, but that supply voltage and current to the power supply through electromagnetic induction. Power transformers often have multiple secondary windings, to increase the isolation of each of a component's multiple power supplies.

selectivity Specification describing an FM tuner's ability to reject unwanted stations. Good selectivity is important to those who live in cities, where stations are closely spaced on the broadcast spectrum.

sensitivity 1) A measure of how much sound a speaker produces for a given amount of input power. Speaker sensitivity is measured by driving a speaker with 1W of power and measuring the sound-pressure level from a distance of 1 meter. 2) A measure of an FM tuner's ability to pull in weak stations. Tuner sensitivity is important if you live a long distance from FM transmitters.

servo An electrical or mechanical system that controls the speed or position of a moving device by forcing it to conform to a desired speed or position. In a servo-controlled turntable, the platter's rotational speed is compared with a reference frequency, and a correction signal is applied to the drive motor to keep the two frequencies identical.

servo-driven woofer A woofer in which an accelerometer (a device that converts motion to an electrical signal) attached to the voice coil sends to an amplifier information about the woofer's position and motion. The amplifier then applies a correction signal so that the woofer's motion more closely matches the characteristics of the audio signal.

shorting plug An audio connector with no cable attached, and whose pins are joined together. When a shorting plug is inserted in an unused preamplifier input, any noise on the unused signal conductor is shorted to ground.

sibilance *s* and *sh* sounds in speech or singing.

signal/noise ratio Numerical value expressing in decibels the difference in level between an audio component's noise floor and a reference signal level.

silky Reproduced midrange or treble sound that lacks hardness or glare.

silver paste A material applied to contact points in an AC power system, such as between circuit breakers and wire, wire and AC outlets, and the point where a house's ground wire is connected to the house's copper water pipe.

single-ended amplifier Amplifier in which both half-cycles of the audio waveform are amplified by the output tubes or transistors. Compare with *push-pull amplifier*.

single-presentation listening Evaluating audio components by listening to only that component, rather than in comparison with other components.

six-channel input A preamplifier input comprising six discrete jacks that will accept the six discrete analog outputs from a multichannel disc player.

skating Force generated in LP playback that pulls the stylus, cartridge, and tonearm toward the center of the record.

skin effect Phenomenon in cables in which high frequencies travel along the conductor's surface, or skin, and low frequencies travel through the conductor's center.

slam Sound with strong, dynamic bass impact that the listener feels in his or her body.

slew rate The speed at which an amplifier can produce an output voltage in response to an input signal. Typically measured in volts per microsecond (V/ μ s).

smooth Sound that lacks peaks and dips in the frequency response.

solid-state drive Storage medium in which data are stored in memory chips rather than on a moving mechanism such as a hard-disk drive.

Sorbothane Rubber-like material that absorbs vibrational energy.

soundfield see *soundstage*

sound-pressure level (SPL) A measure of loudness expressed in decibels (dB).

soundstage A recording's reproduction of an apparent three-dimensional soundspace in front of and/or around the listener.

source component Component that provides audio and/or video signals to the rest of an audio system; e.g., DACs, turntables, music servers, FM tuners. Compare with *control components* and *playback components*.

source switching Function performed via a preamplifier or integrated amplifier, in which the user selects which source component's signals are fed to the speakers.

spade lug A speaker termination with a flat area that fits around a binding post.

S/PDIF (Sony/Philips Digital Interface Format) interface Standardized method of transmitting digital audio between components.

speaker see *loudspeaker*

SPL see *sound-pressure level*

SPL meter Device for measuring the sound-pressure level created by an audio source.

sprung turntable A turntable in which the platter and armboard are mounted on a subchassis that floats on springs within the base. Compare with *unsprung turntable*.

SQ A Quadraphonic format invented by CBS (later bought by Sony) in which the rear channels were matrix-encoded into the front channels.

SSD see *solid-state drive*

standing wave Stationary areas of high and low acoustical pressure in a room; caused by interaction of the soundwaves with the room's boundaries.

stator The electrostatically charged element in an electrostatic loudspeaker driver that remains stationary.

step-up transformer Transformer with a higher output voltage than the voltage input to it. Sometimes used between a moving-coil cartridge and a moving-magnet phono input.

stopband Range of frequencies a filter does not let through or pass along.

strain-gauge cartridge Type of phono cartridge that does not operate by electromagnetic induction, as do moving-magnet and moving-coil cartridges. A specialized preamplifier delivers a current through semiconductor crystals in the cartridge that change their resistance as a result of mechanical pressure caused by

modulations in the groove. The varying resistance creates a varying current that is decoded as an audio signal.

streaming Delivering audio content to a digital-to-analog converter for real-time listening rather than to store the content.

ST-Type (Straight Tip) optical interface Transmission system for sending digital audio in the S/PDIF format in which the signal is transmitted as light down a glass conductor. An ST-Type optical cable is terminated with locking bayonet connectors. No longer used on modern audio equipment.

stylus Tiny wedge protruding from a phono cartridge's cantilever that rests in the record groove and moves back and forth in response to modulations in the groove.

stylus cleaner Fluid that removes dirt and debris from a cartridge stylus.

stylus-force gauge An instrument that measures the amount of downward force a stylus exerts on the record; aka a tracking-force gauge.

subchassis In a turntable, the platform beneath the plinth (base) that supports the platter bearing. The subchassis is often suspended on springs to isolate it from vibration.

subscription-based music service Method of accessing recordings in which you pay a monthly fee to stream recordings from the Internet.

subwoofer A loudspeaker designed to reproduce only low-bass frequencies.

Super Audio CD (SACD) Format developed by Sony and Philips for storing CD-quality audio as well as high-resolution digital audio, and sometimes multichannel audio, on a CD-sized disc. Hybrid SACD/CDs can be played in conventional CD players as well as in SACD players.

supra-aural (literally, "above the ear") Headphones in which the transducers rest directly on the ears, aka "on-ear" headphones. Compare with *circumaural headphones*, which go *around* the ears.

surround decoder A circuit or component that converts a surround-encoded audio signal into multiple audio signals that can be amplified. All A/V controllers incorporate surround decoding to convert the Dolby Digital or DTS signals from DVD, satellite, BD, or over-the-air HDTV broadcasts to discrete analog signals.

surround decorrelation A THX technology that makes the sound in the monaural left and right surround channels slightly different, thus increasing the senses of envelopment, diffusion, and space.

surround delay A technique of delaying the signal to the surround channels to increase the apparent separation between the front and surround channels and thus create a more convincing sense of space.

surround mode A setting on A/V receivers and A/V preamplifiers that determines what, if any, surround decoding or digital signal processing is performed on the audio signal.

surround receiver see *A/V receiver*

surround sound Audio recording and playback formats that use more than two channels and are reproduced using more than two loudspeakers, some of which are located behind the listener.

surround speakers Speakers located beside or behind the listener that reproduce the surround channels of surround-sound-encoded audio programs.

S-video A video connection method that keeps separate the video signal's brightness and color information.

sweet spot Point in the listening room where the sound is best.

system matching Combining components to create the most musical-sounding system for a given budget.

tag Non-audio information attached to a music file, aka *metadata*.

tagging Manually adding *metadata (tags)* to a music file.

tangential-tracking tonearm A tonearm in which the entire tonearm and tonearm bearing move in relation to the record along a fixed tangent so that the relationship between stylus and groove remains constant throughout the entire LP. Compare with the more common *pivoted tonearm*, in which the arm, cartridge, and stylus traverse an arc across the record, which means that the angle or tangent of stylus to groove is constantly changing.

taut Tight, well-defined reproduction of bass frequencies.

terminated cable A cable employing resistors, inductors, and capacitors in a box that is attached to and part of the cable.

terminations The fittings on the end of a cable: RCA plugs, spade lugs, banana plugs, XLRs, etc.

THD see *total harmonic distortion*

theater pass-through Feature on some preamplifiers and integrated amplifiers that sets the preamplifier to some fixed gain setting (usually unity gain) independent of the volume-control setting. Useful in systems that also include a digital controller; the theater pass-through input allows the controller to maintain the same gain relationships among all channels, regardless of the volume setting on the preamplifier or integrated amplifier.

thick 1) Bass with poor pitch definition; 2) an opaque soundstage.

thin Lacking bass weight, substance, heft.

three-way speaker A loudspeaker that divides the audioband into three parts (bass, midrange, treble) for reproduction by three or more drivers.

THX Lucasfilm's set of patents, technologies, and technical/acoustic performance criteria for filmsound reproduction in movie theaters. See also *Home THX*.

THX-certified An A/V product that correctly implements the THX technologies and meets stringent technical performance criteria for film-sound reproduction.

THX EX Surround format that matrix-encodes a third surround channel in the existing left and right surround channels of a Dolby Digital signal. This third surround channel drives a loudspeaker(s) located directly behind the listening position for a more convincing sense of space. Called Dolby Digital EX in non-THX-certified products.

timbre The quality of a tone produced by a musical instrument or voice.

timbre matching A THX technology that ensures that sounds arriving from the sides have the same timbres as sounds arriving from in front of the listener, to ensure the smooth panning of sounds.

time smearing Phenomenon that occurs in digital filters in which some energy in the audio signal is spread out over time, before and after that signal.

tizzy In sound, excessive upper-treble energy that is also grainy.

toe in Angling loudspeakers to point directly toward the listening position rather than straight ahead.

tonal balance Relative levels of bass, midrange, and treble in the sound of an audio component or recording.

tonearm Component of an LP playback system that holds the phono cartridge.

tonearm effective length Distance between the tonearm's pivot point and the stylus.

TosLink cable Optical cable for carrying digital audio developed by Toshiba (from Toshiba Link).

total harmonic distortion (THD) Specification stating the amount of harmonic distortion in an audio component. Called "total" because it is the sum of all the individual harmonic-distortion components created by the component.

tracking error Difference between the angle of the cutting stylus to an LP groove and the angle of the playback stylus.

tracking-force gauge see *stylus-force gauge*

transducer Any device that converts energy from one form to another. Microphones, loudspeakers, and phono cartridges are transducers. A microphone converts acoustic energy to electrical energy; a loudspeaker converts electrical energy to mechanical energy.

transient A sound that begins very suddenly, often at high level; e.g., the sound of a snare drum struck sharply with a hard drumstick.

transient response The ability of a component to accurately reproduce sudden sharp sounds.

transistor Device made from solid semiconductor material that can amplify audio signals.

transition band Range of frequencies in a filter between the passband (the frequency range passed by the filter) and the stopband (frequency range not passed by the filter).

transmission line A loudspeaker loading technique in which the woofer's rear wave is absorbed by a long duct or series of baffles to create a labyrinth (the transmission line) inside the enclosure.

transparent Sound of a component or system with very low levels of coloration. A soundstage in which the reproduction of acoustic space sounds clear rather than veiled.

treble High audio frequencies; generally, the range from 3kHz to 20kHz.

Trifield A pair of algorithms derived from Ambisonics in which a center-channel signal—and, in some applications—rear-channel signals, are derived from a 2-channel recording.

triode The simplest vacuum tube, employing just three elements: the cathode, plate, and control grid.

TRRS (tip, ring, ring, sleeve) plug Four-conductor plug most often found on balanced headphones and balanced headphone amplifiers. The four conductors are left-channel +, left-channel –, right-channel +, and right-channel –. TRRS plugs are also used in some smartphones to connect unbalanced earbuds plus a microphone signal.

TRS (tip, ring, sleeve) plug Headphones plug that carries two signals (left- and right-channel audio) plus a ground on the plug's three conductors: the tip, ring, and sleeve. TRS plugs can be 2.5mm, 3.5mm, or ¼" in diameter.

tube see *vacuum tube*

tube damper Device that fits over a vacuum tube to reduce vibration. Some tube dampers also act as heatsinks.

turnkey music server A music server that requires no user setup or configuration. Compare with computer-based *do-it-yourself music server*.

tweeter A speaker driver designed to reproduce only high-frequency (treble) signals.

two-way speaker A loudspeaker that splits the audioband into bass and treble bandwidths for reproduction by two or more drivers.

UHJ In Ambisonics, the mono-and-stereo-compatible combination of the W (mono pressure signal), X (sides), and Y (front to back) signals.

ultralinear A method of operating a pentode vacuum tube that combines the output power of a pentode tube with the distortion characteristics of a triode tube. Invented by David Hafler and Herbert Keroes in 1951.

ultrasonic record cleaner A cleaner in which the record is immersed in liquid, through which ultrasonic energy dislodges dirt and debris.

unbalanced connection Connection method in which the audio signal is carried on two conductors, the signal and ground. Compare with *balanced connection*, in

which the audio signal is carried on three conductors, two signal conductors and ground.

unbalanced headphone drive Method of powering headphones in which an amplifier is connected to one side of the driver and the other side is connected to a common ground. Compare with *balanced headphone drive*, in which two amplifiers, each connected to one side of the driver, power the headphones.

underdamped Describes a loudspeaker whose woofer(s) continues to move after the drive signal has stopped.

unipivot Type of tonearm bearing in which the arm rests on and pivots around a single point. Compare with *gimbal bearing*.

unipolar loudspeaker A loudspeaker that directs its acoustic output in one direction. Compare with *dipolar loudspeaker*.

universal disc player Source component that plays CD, DVD-Video, DVD-Audio, Blu-ray, and SACD.

Universal Plug and Play (UPnP) Communication protocols in consumer-electronics devices that allow them to connect to a network and recognize other devices on the network without user intervention.

Universal Serial Bus (USB) Computer interface format that has been adapted to transmit digital audio data. See also *adaptive USB* and *asynchronous USB*.

unsprung turntable Turntable in which the platter and tonearm are connected directly to the turntable base rather than suspended from a subchassis. Compare with *sprung turntable*.

upconvert see *upsample*

UPnP see *Universal Plug and Play*

upsample Technique of increasing the sampling frequency of a digital audio signal so that it can be converted to analog from the higher frequency.

uptilted A sound having too much treble and not enough bass.

usable sensitivity Tuner specification that states the voltage across the antenna required to produce an audio signal with a signal/noise ratio of 30dB. Compare with the more stringent *quieting sensitivity*.

USB see *Universal Serial Bus*

user interface The controls and displays on a product and their ease of use.

vacuum hold-down Mechanism for achieving tight coupling between an LP and a turntable platter by creating a vacuum between LP and platter.

vacuum tube Device for amplifying audio in which the active elements are enclosed in a glass envelope devoid of air.

value product Audio component designed to achieve maximum sonic performance for the lowest price.

veiled Impression of a scrim or haze between the listener and the sound. Compare with *transparent*.

vented loudspeaker see *reflex loudspeaker*

vertical tracking angle (VTA) The angle at which the stylus sits in a record's groove. Adjusted by moving the tonearm pivot point up or down.

vertical tracking force (VTF) Pressure applied to the stylus in a record groove.

vertical venetian-blind effect Change in perceived tonal balance when the listener moves sideways in the listening position. Generally occurs only with electrostatic loudspeakers because of their narrow dispersion.

video display Device that converts a video signal into a visual image.

vivid Reproduced sound in which every instrument and voice is clearly audible. An overly vivid presentation is *analytical* and unmusical.

voice coil Loop of wire wrapped around a bobbin positioned in the center of a loudspeaker driver through which flows current from the power amplifier.

volt (V) Unit of electromotive force. The difference in potential required to make one ampere of current flow through one ohm of resistance. See also *voltage*.

voltage Analogous to electrical pressure. Voltage exists between two points when one point has an excess of electrons in relation to the other point. A good example is a battery: its negative terminal has an excess of electrons in relation to its positive terminal. If you connect a piece of wire between a battery's positive and negative terminals, voltage pushes current through the wire. One volt across 1 ohm of resistance produces a current of 1 ampere.

voltage source A power amplifier that can maintain its output voltage regardless of the load impedance.

VTA see *vertical tracking angle*

VTF see *vertical tracking force*

wall-stud resonance Acoustic resonance of a listening room's walls when struck by sound. Occurs only with drywall-on-stud construction.

WASAPI (Windows Audio Session Application Programming Interface) Software for Windows computers that manages the flow of audio data between an application and the audio device connected to the computer. Allows an uncorrupted bitstream to be delivered by the computer to the audio device. Similar to *ASIO*.

watermarking 1) Embedded digital codes disguised as noise in the audio waveform to control how content is copied, and to trace the source of music piracy. 2) A visible imprint on the surface of a legitimate SACD that disappears on a copy.

watt W Unit of electrical power defined as the power dissipated by 1 ampere of current flowing through 1 ohm of resistance.

WAV (Waveform Audio File) Digital audio file format in which the data are stored or transmitted in raw PCM format.

wavelength The distance between successive cycles of a sinewave or other periodic motion.

weighting curve A filter applied to noise or sound-pressure-level measurements that approximates the perceived loudness or noise level so that the measurement better correlates with subjective experience.

weighty Full, robust reproduction of bass.

WiFi Communication protocols that allow devices to wirelessly communicate with each other or connect to the Internet.

Windows Audio Session API see *WASAPI*

Windows Media Audio Lossless (WMA Lossless) File format that provides about 40% reduction in file size with perfect bit-for-bit accuracy to the source data.

woofer Driver in a loudspeaker that reproduces bass frequencies.

XLR jack and plug Three-pin connector that usually carries a balanced audio signal.

zero-crossing distortion Discontinuity in a musical waveform in a push-pull amplifier.