MIKE SENIOR

MIXING SECRETS

FOR THE SMALL STUDIO

Focal Press

Mixing Secrets for the Small Studio

Mike Senior

Amsterdam • Boston • Heidelberg • London • New York • Oxford • Paris
San Diego • San Francisco • Singapore • Sydney • Tokyo
Focal Press is an imprint of Elsevier

FocalPress is an imprint of Elsevier
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
The Boulevard, Langford Lane, Kidlington, Oxford, OX5 1GB, UK

Notices
Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data
Senior, Mike.
Mixing secrets for the small studio / Mike Senior.
p. cm.
ISBN 978-0-240-81580-0 (pbk.)
1. Sound—Recording and reproducing. 2. High-fidelity sound systems. 3. Sound studios. I. Title.
TK7881.4.S465 2011
621.389'3—dc22 2010045009

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library.

ISBN: 978-0-240-81580-0

For information on all Focal Press publications
visit our website at www.elsevierdirect.com

11 12 13 14 15 5 4 3 2 1

Printed in the United States of America

To my parents.

This page intentionally left blank

Contents

This page intentionally left blank

Acknowledgments

This book has been in the works a long time, and many people have lent their support during its creation. First of all, I'd like to thank all those who peppered me with questions at the Cambridge Music Technology masterclass sessions, thereby sowing the seeds of this book, and also the many readers of *Sound on Sound* magazine whose numerous "Mix Rescue" submissions and questions have clarified my thinking.

In addition, I'd like specifically to thank all the interviewers who have done an immense service to us all by shedding so much light on top-level studio practice: Michael Barbiero, Matt Bell, Bill Bruce, Richard Buskin, Dan Daley, Tom Doyle, Maureen Droney, Tom Flint, Keith Hatschek, Sam Inglis, Dave Lockwood, Howard Massey, Bobby Owsinski, Andrea Robinson, and Paul Tingen. Paul Tingen deserves special praise for his dogged pursuit of the hottest current hitmakers for *Sound on Sound*'s "Inside Track" series. I'm grateful as well to Russ Elevado, Roey Izhaki, Roger Nichols, and Mike Stavrou for their own insightful writing on the subject of mixdown. Many thanks are also owed to Philip Newell, Keith Holland, and Julius Newell for permission to reproduce the results of their superb NS10M research paper; to Phil Ward for alerting me to the perils of speaker porting; and to Roberto Détrée and Mastermix Studios in Munich for allowing me to photograph their speakers.

In developing this text for publication, I have been assisted a great deal by Matt Houghton and Geoff Smith, whose well-informed and in-depth feedback has been invaluable. Thanks also to everyone in the editorial department at *Sound on Sound* for generously offering so much help and useful advice. I'm also very grateful to the team at Focal Press for their patience and expertise in bringing this project to fruition: Catharine Steers, Carlin Reagan, Melissa Sandford, Laura Aberle, and Graham Smith.

Above all, I'd like to thank my wonderful wife, Ute, for her unwavering love and support, as well as for taking on the worst proofreading and referencing tasks so graciously. I'd be lost without you, my love. And thank you Lotte and Lara too—yes, Papa's finished being boring now.

This page intentionally left blank

Introduction

WHAT YOU'LL LEARN FROM THIS BOOK

This book will teach you how to achieve release-quality mixes on a budget within a typical small-studio environment by applying power-user techniques from the world's most successful producers. Using these same methods, I've carried out dozens of mix makeovers for *Sound on Sound* magazine's popular "Mix Rescue" series, working on mass-market gear in various home, project, and college studios. If you head over to *www.soundonsound.com*, you can find before/after audio comparisons for every one of these remixes, and this book is a one-stop guide to pulling off these kinds of night-and-day transformations for yourself.

WHAT YOU WON'T LEARN

This book will not teach you how to operate any specific brand of studio gear—that's what equipment manuals are for! The information here is deliberately "platform neutral," so that you can make just as much use of it whether you're on Cubase, Digital Performer, Live, Logic, Pro Tools, Reaper, Reason, Sonar, or any other software platform. And although I've made the assumption that the majority of cost-conscious mix engineers will now be working in software, my advice also applies equally well to hardware setups, give or take a patch cord or two. Indeed, my own background is in computerless environments, so I know from experience that equally good results are attainable there.

WHAT YOU NEED TO KNOW ALREADY

Although I've done my best to make this book friendly to studio newbies, there is nonetheless some basic background knowledge that you'll need to understand to get the best out of what I'll be writing about. In particular, I'm assuming that the reader already understands something about these topics:

- *The fundamental physics, measurement, and perception of sound*: amplitude, decibels, and loudness; frequency, Hertz, and pitch; sine waves and the harmonic series; frequency response measurements
- *Studio setup and session workflow*: transmission/storage methods for sound (electrical, magnetic, digital); the basic principles of mics, DIs, audio cables, multitrack recorders, and mixers; routing for insert and loop effects; the stages of a typical production (preproduction, tracking, overdubbing, mixing, mastering); gain, headroom, noise, and signal metering

If you need a quick refresher on any of these elements, then check out *www.cambridge-mt.com/ms-basics.htm* for a whistle-stop overview. Alternatively, you can get more into the equations by checking out the first two chapters of either Alexander U. Case's *Sound FX* (Focal Press, 2007) or David Miles Huber & Robert E. Runstein's *Modern Recording Techniques*, 7th edition (Focal Press, 2009).

Studio jargon can be pretty intimidating, but the sooner you get a grip on it, the quicker you'll improve your mixing. If you feel unsure of any of the terminology used in this book, then head over to *www.cambridge-mt.com/ms-links.htm*, where there are links to a couple of thorough and well-maintained glossaries.

HOW TO USE THIS BOOK

Because this book has been specifically designed as a step-by-step primer, you'll get best results if you work through it from beginning to end. Many later sections rely on material covered in earlier chapters, so some aspects of the discussion may not make the greatest sense if you just dip in and out. At the end of each chapter there is a **Cut to the Chase** section, which allows you to review a summary of each chapter's main "secrets" before proceeding. Underneath it is an **Assignment** section, which suggests a number of practical activities to consolidate your understanding, and these assignments could also serve as coursework tasks within a more formal education framework. The URL at the end of each chapter leads to a separate website containing a selection of related links and audio files, all of which may be freely used for educational purposes.

This book is based on my own extensive research into the studio practices of more than 100 world-famous engineers, drawing on more than 2 million words of firsthand interviews. The text therefore includes hundreds of quotes from these high-fliers. If you don't recognize someone's name, then look it up in **Appendix 1** to get an idea of the most high-profile records they've worked on—you'll almost certainly have heard a few of those! If you'd like to read any quote in its original context (which I'd heartily recommend), then follow the little superscript number alongside it to **Appendix 2,** where there's full reference information for each one. Finally, if you have any further questions or feedback, feel free to email me at **ms@cambridge-mt.com**.

PART 1
Hearing and Listening

Probably the most reliable way to waste your time in a small studio is by trying to mix before you can actually hear what you're doing. Without dependable information about what's happening to your audio, you're basically flying blind, and that can get messy. In the first instance, you'll face a frustratingly uphill struggle to get a mix that sounds good in your own studio, and then you'll invariably find that some of your hard-won mixes simply collapse on other playback systems, so that you're left unsure whether any of the techniques you've learned along the way are actually worth a brass farthing. You'll be back to square one, but with less hair.

Relevant advice from professional engineers is perhaps unsurprisingly thin on the ground here. After all, most pros have regular access to expensive high-end speaker systems in purpose-designed rooms with specialist acoustic treatment. However, even the hottest names in the industry don't always get to work in the glitziest of surroundings, and if you look carefully at their working methods, they have actually developed various tactics that enable them to maintain consistent high-quality results even under difficult circumstances. These same tricks can be applied effectively in small studios too. So much so, in fact, that as long as you take care with gear choice and studio setup, it's perfectly possible to produce commercially competitive mixes in a domestic environment with comparatively affordable equipment. Indeed, all of my remixes for *Sound on Sound* magazine's monthly "Mix Rescue" column have been carried out under exactly such restrictions.

But even God's own personal control room won't help you mix your way out of a wet paper bag unless you know how to *listen* to what you're hearing. In other words, once you're presented with a bunch of information about your mix, you need to know how to make objective decisions about that data, irrespective of your own subjective preferences, because that's the only way of repeatedly meeting the demands of different clients or different sectors of the music market. Do the cymbals need EQ at 12kHz? Does the snare need compression? How loud should the vocal be, and are the lyrics coming through clearly enough? These are the kinds of important mix questions that neither your listening system nor your mixing gear can answer—it's you, the engineer, who has to listen to the raw audio facts, develop a clear opinion about what needs to be changed, and then coax the desired improvements out of whatever equipment you happen to have at your disposal.

Most people who approach me because they're unhappy with their mixes think that it's their processing techniques that are letting them down, but in my experience the real root of their problems is usually either that they're not able to hear what they need to, or else that they haven't worked out how to listen to what they're hearing. So instead of kicking off this book by leaping headlong into a treatise on EQ, compression, or some other related topic, I want to begin instead by focusing on hearing and listening. Until you get a proper grip on those issues, any discussion of mixing techniques is about as useful as a chocolate heatsink.

CHAPTER 1

Using Nearfield Monitors

1.1 CHOOSING YOUR WEAPONS

Choosing the equipment that allows you to hear (or "monitor") your mix signal is not a task to be taken lightly, because it's the window through which you'll be viewing everything you do. For those on a strict budget, however, the unappetizing reality is that monitoring is one of those areas of audio technology where the amount of cash you're prepared to splash really makes a difference. This is particularly true with regard to your studio's primary monitoring system, which needs to combine warts-and-all mix detail with a fairly even frequency response across the biggest possible slice of the 20Hz to 20kHz audible frequency spectrum—a set of characteristics that doesn't come cheap.

That said, when choosing the stereo loudspeakers that will fulfill these duties in all but the most constrained studios, there's a lot you can do to maximize your value for money. First off, furniture-rattling volume levels aren't tremendously important for mixing purposes, despite what you might guess from seeing pics of the dishwasher-sized beasts mounted into the walls of famous control rooms—most mix engineers use those speakers mainly for parting the visiting A&R guy's hair! "I use nearfields almost exclusively," says Chuck Ainlay, "because there just aren't many situations where the main monitors sound all that good. The mains in most studios are intended primarily for hyping the clients and playing real loud."[1] "I don't use the big monitors in studios for anything," says Nigel Godrich, "because they don't really relate to anything."[2] You'll get a more revealing studio tool at a given price point if you go for something where the designers have spent their budget on audio quality rather than sheer power. As it happens, the most high-profile mix engineers actually rely almost exclusively on smaller speakers set up within a couple of meters of their mix position (commonly referred to as nearfield monitors). If you sensibly follow their example in your own studio, you shouldn't need gargantuan speaker cones and rocket-powered amplifiers, even if you fancy making your ears water.

Mixing Secrets for the Small Studio.

SURROUND MONITORING

Before acquiring a multispeaker surround setup for a small studio, I'd advise thinking it through pretty carefully. Until you can reliably get a great stereo mix, I for one see little point in spending a lot of extra money complicating that learning process. In my experience, a limited budget is much better spent achieving commercial-quality stereo than second-rate surround, so I make no apologies for leaving the topic of surround mixing well alone and concentrating instead on issues that are more directly relevant to most small-studio denizens.

Another simple rule of thumb is to be wary of hi-fi speakers, because the purpose of most hi-fi equipment is to make everything sound delicious, regardless of whether it actually is. This kind of unearned flattery is the last thing you need when you're trying to isolate and troubleshoot sneaky sonic problems. I'm not trying to say that all such designs are inevitably problematic in the studio, but most modern hi-fi models I've heard are just too tonally hyped to be of much use, and maintenance issues are often a concern with more suitable pre-1990s systems. Speakers with built-in amplification (usually referred to as "active" or "powered') are also a sensible bet for the home studio: they're more convenient and compact; they take the guesswork out of matching the amplifier to your model of speaker; they're normally heavier, which increases the inertia of the cabinet in response to woofer excursions; and many such designs achieve performance improvements by virtue of having separate matched amplifiers for each of the speaker's individual driver units.

Beyond those issues, a lot of monitor choice is about personal preference, and there's nothing wrong with that. Some people prefer bright aggressive-sounding monitors, others restrained and understated ones, and neither choice is wrong as such. The main thing to remember is that no monitors are truly "neutral," and every professional engineer you ask will have his or her own personal taste in this department. Part of the job of learning to mix is getting accustomed to the way your particular speakers sound, so don't get too uptight about minute differences in tone between speakers. Go for something that appeals to you, and then concentrate on tuning your ears to how your chosen model responds in your own control room. "You've got to be careful about getting new monitors," advises Dave Way. "You've got to break them in and get to know them before you start to rely on them."[3] Part of doing this involves referring to a set of reference recordings with which you're familiar (discussed more in Chapter 4).

No monitors are truly "neutral," and every professional engineer you ask will have his or her own personal taste in this department. Part of the job of learning to mix is getting accustomed to the way your own particular speakers sound.

Ported Speakers and Frequency Response

I have one further piece of advice to offer when choosing monitors, but I've deliberately held it in reserve, because I want to give it special attention. It's this: the less money you have to spend, the more you should beware ported monitors. Such speakers are sometimes also referred to as "bass reflex" or "reflex loaded" designs, and they incorporate some kind of hole or vent in the speaker cabinet, which encourages the whole box to resonate in sympathy with the speaker's drivers. The main purpose of this resonance is to increase the low-frequency output, an aspect of a small speaker's performance that is naturally restricted based on its limited woofer size. By using a port to compensate for the woofer's natural low-end roll-off, manufacturers can have a wider flat region on their published frequency-response graph, as well as give the speaker a louder, beefier sound that'll help impress Joe Public's wallet in the shops. Figure 1.1 illustrates the basic effect of porting on a typical small-studio monitor's low-end frequency response. The solid line on the graph shows the kind of response you'd expect of a fairly typical small ported speaker, with the output remaining within a ±3dB window down to maybe 55Hz. If you defeated the speaker's port by blocking it, however, you'd find that the response changed to something like that shown by the dotted line: the trace now drifts out of the ±3dB window almost an octave higher, just above 100Hz.

So what's so bad about using a port to widen a speaker's frequency response? The problem is that porting also has several less well-advertised side effects that

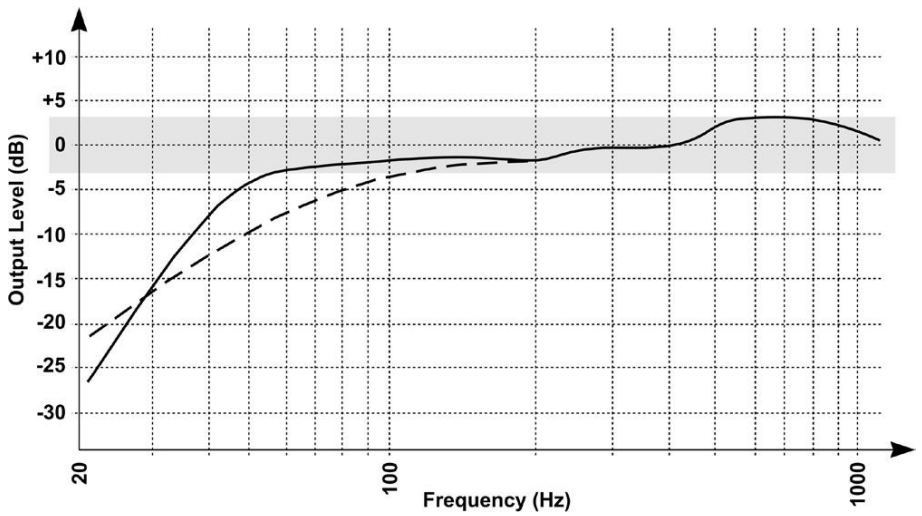

FIGURE 1.1

The solid line on this graph shows the kind of sub-1kHz frequency response plot you might expect for a small and budget-friendly ported studio monitor. The dotted line shows how the response changes when the port is artificially blocked, thereby defeating the cabinet's low-frequency resonance. The shaded region indicates the ±3dB "flat" region of the speaker's quoted frequency-response specification.

can easily conspire to hamstring you at mixdown. Given the widespread use of porting in budget nearfield monitors, it's important to understand what these side effects of porting are. On the one hand, this knowledge makes it easier to evaluate monitors objectively when making a purchase; on the other hand, it better equips you to work around potential porting gremlins when the choice of monitors is beyond your control—for example, in a college facility or a friend's home studio. So bear with me while I look at this issue in more detail.

The first problem with porting can already be seen in Figure 1.1: although the port stops the response dropping off until 50Hz, the output takes a real nosedive beyond that. This means that although the speaker's overall low-frequency output is boosted by the port, the relationship between the sub-50Hz levels and the rest of the signal is seriously skewed at the same time, which makes it trickier to make judgments about instruments with important low-frequency components. So assuming, for the sake of example, that you're playing back the sound of a bass instrument that is completely consistent in its low-frequency levels, the perceived volume of its fundamental frequency will still dance around alarmingly as the notes change pitch, depending on how far the fundamental slips down the steep frequency-response roll-off.

Bear in mind that the lowest fundamental from a bass guitar is around 44Hz, whereas pianos, organs, and synths are just some of the sources that will happily generate fundamentals in the 20 to 40Hz bottom octave. In contrast to the fundamental, however, the first harmonic of these bass notes lies an octave above, typically in the much flatter frequency-response region above 50Hz, so it'll be tough going to decide whether there's the right amount of each of these frequencies respectively. And, of course, if we step back into the wild again, where untamed rampaging bass parts are often anything but consistent, how are you expected to judge when your mix processing has actually reined them in properly?

Kick drums are equally complicated to deal with. Let's say that you're comparing the kick level in your own mix to something on a favorite commercial record, but your kick drum has loads of energy at 30Hz, whereas the comparison track's kick is rich in the 50Hz region. Because the speaker is effectively recessing the 30Hz region by 12dB compared to the 50Hz region, you're likely to fade your own kick drum up too high, only to discover a rumbling mess lurking underneath your mix in other monitoring environments. Although the loss of low end on an unported monitor is also a problem, it's much easier to compensate for this mentally while mixing, because the relative levels of neighboring low-frequency bands are more representative.

Killer Side Effects of Porting

These porting anomalies, however, are only the tip of the iceberg, because frequency-response graphs only show how speakers respond to constant full-frequency noise, a test signal that is nothing like the varied and fast-moving waveforms of music. Much more troublesome is the way that porting hinders

the monitor's ability to track moment-to-moment changes in the mix signal. Specifically, the port causes any spectral energy at its resonant frequency to ring on for a short time, and while it's this resonant buildup that generates the port's flattering low-frequency level boost for a constant noise test signal, the same quality also adds short resonant tails to fleeting percussive attack noises (often referred to as transients), such that they can seem louder and less punchy than they actually are. Sounds that stop abruptly suffer a similar problem, with the port ringing on after they've finished. In this case, the resonance not only disguises the true decay attributes of the sound itself, but it can also make it difficult to judge the character and level of short-duration studio effects (such as modulated delays and reverb), which are often very useful at mixdown.

Another possible problem with ported speakers is that the ringing of the port can dominate over the real fundamental frequencies of low bass notes, making them difficult to distinguish from each other. Speaker reviewers sometimes refer to this phenomenon as "one-note bass," and it adds unwelcome uncertainty to tuning judgments at the low end. A commercial recording that I find particularly good for revealing this occurrence is Skunk Anansie's "Infidelity" (from the band's album *Stoosh*), where the meandering bass line quickly becomes murky and ill-defined in the presence of low-end monitoring resonances. (The track is also good for testing the frequency response of a monitoring system, as only the most extended response can do justice to that particular kick drum's almost seismic low-frequency rumble.)

Were the port-ringing consistent across the audio spectrum, you could mentally compensate for it perhaps, but of course it's not: it's more or less severe depending on how much of a given transient's energy resides around the porting frequency. Furthermore, I've so far taken for granted that the port has only one resonant frequency. In reality, however, it's difficult to stop the thing resonating at a whole range of higher frequencies too, which leads to unpredictable time-smearing artifacts right across the frequency spectrum. So it's not just bass instruments that you may be unable to judge reliably, but everything else too! Although it's perfectly possible for speaker designers to use careful internal cabinet design and damping to tame all but the desired low-frequency port resonance, that does cost them money, so this is where more affordable designs can really come a cropper.

Of course, a simple frequency-response graph leaves you blissfully ignorant of any of this stuff, because it only has axes for frequency and level. If you want to lay bare resonance side effects, then you need to add a third dimension to your frequency-response graph: time. Fortunately, there is a type of graph that does exactly that, called a spectral decay or "waterfall" plot. It reveals what happens to a speaker's output when a constant full-range test signal is suddenly switched off—as the graph develops in time (in other words moving from the background into the foreground, speaking three-dimensionally), you can see how much different frequencies ring on.

FIGURE 1.2
Some affordable two-way ported nearfield monitors (*top to bottom*): the ADAM A7X has dual circular ports either side of the woofer, the KRK Rokit 8 has a port slot under the woofer, the Behringer 2030A has two narrow port slots on either side of the tweeter, and the M-Audio DSM2 features a single circular port on the rear panel.

The left-hand column of Figure 1.3 shows waterfall plots for three well-designed small nearfield monitors. The top graph is for an unported model, whereas the two lower plots are for ported designs. You can see the low end of the ported models ringing on, as you'd expect, but otherwise the midrange and high end stop quickly without any obvious resonant trails. Compare this with the waterfall plots in the right-hand column of Figure 1.3, measured from three budget ported nearfields. What they clearly illustrate is that each speaker has prominent resonances well up into the midrange. (Bear in mind, though, that other factors contribute unwanted midrange resonances to speakers as

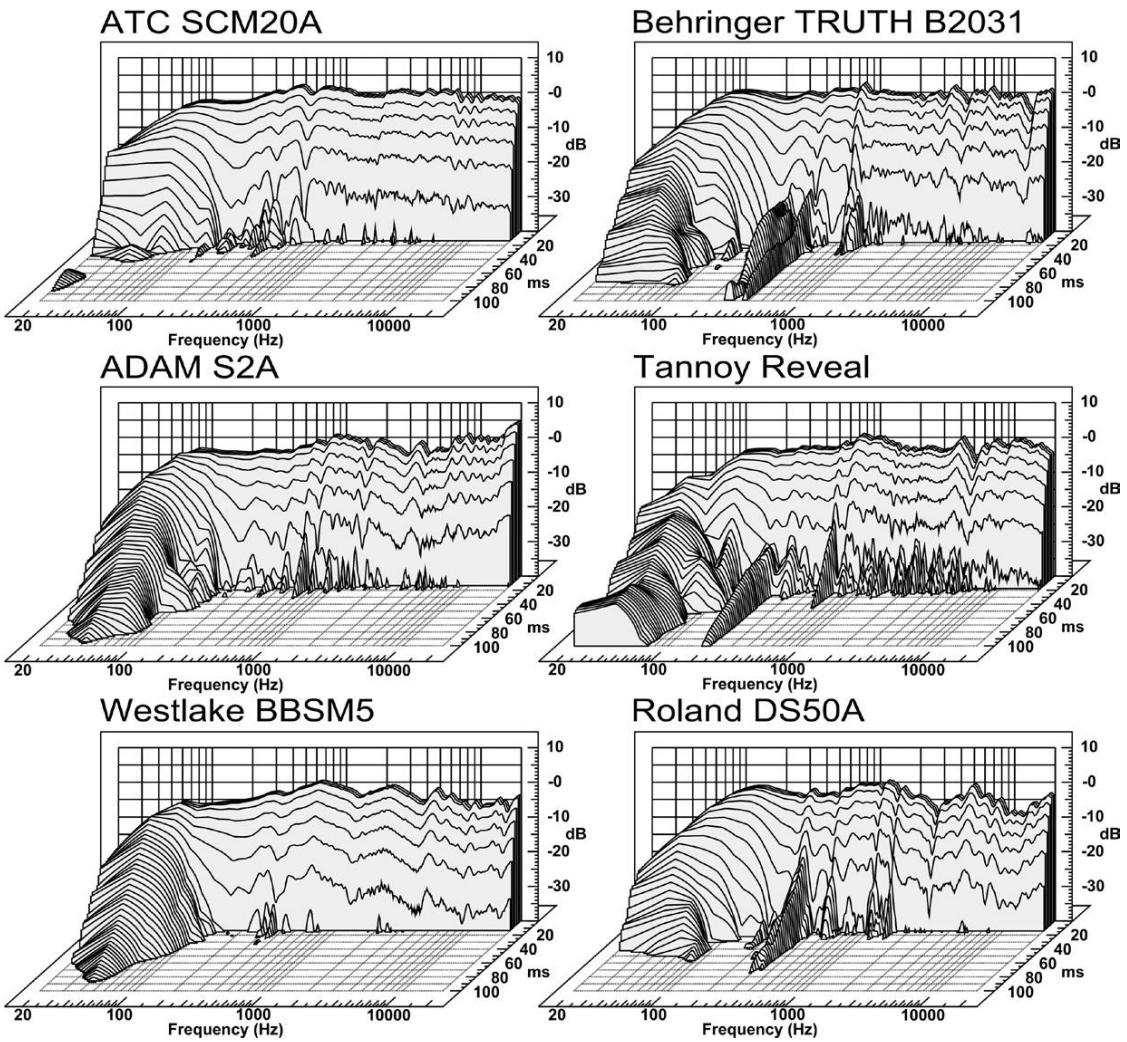

FIGURE 1.3

Waterfall plots for six different sets of studio monitors.

well, so it's not just porting side effects you're seeing here.) The less well-controlled a monitor's resonances, the less easily you can mix with it.

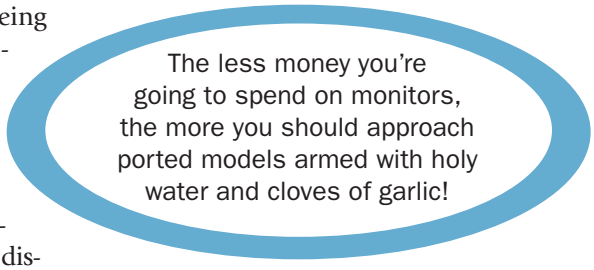

But even that's not the end of the story: ports can also produce turbulence noise, which obscures other parts of your mix; compression artifacts, which mess with the apparent level of bass instruments as you adjust the monitoring volume; and distortion, which misleadingly gives extra midrange body to bass instruments, making them seem more audible in the mix than they should be. If you want to hear what I'm talking about, try listening to the low-frequency sine-wave tones in the LFSineTones audio file through a budget ported monitor. (You can download this file from this chapter's URL: *www.cambridge-mt.com/ms-ch1.htm*.) Particularly on the lowest frequencies you'll usually hear a good dose of fluttering port noise and low-level distortion harmonics overlaid on what should be pure tones. Need any further convincing? Then consider the fact that two of the most influential mixing speakers in the history of audio production are unported designs: the Yamaha NS10 and the Auratone 5C Super Sound Cube. (You can see the waterfall plots for these speakers in Figure 1.5, and although neither has a particularly flat frequency response, both are extraordinarily well-behaved as far as resonances are concerned.)

All of which brings me back to my main point: the less money you're going to spend on monitors, the more you should approach ported models armed with holy water and cloves of garlic! In my experience, you'll have to part with well over £1000 ($1500) for a pair of ported nearfields that can reliably deliver what you need to mix competitively, whereas I don't think you need to spend this much on an unported design to get similar mixing muscle, just so long as you're willing to work with lower overall volume levels. This is why I usually point cash-strapped would-be mix engineers in the direction of products such as NHT Pro's M-00 and S-00 combination, or Blue Sky's Media Desk and Pro Desk systems, all of which are active, unported, and fairly full range.

Speaker Stands and Other Mounting Hardware

You can fork out for the fanciest monitors you like, but unless you set them up sensibly in your room, you might as well have spent most of that money on doughnuts for all the good it'll do your sound. I've visited a large number of small studios, and one thing the majority have in common is that their owners have underestimated the importance of monitor installation, with the result that the monitoring sounds only a fraction as good as it should, given the cost of the speakers. So let's look at ways you can maximize the quality of the sound, whatever speakers you're using.

For a start, the speaker cabinets should be as firmly fixed as possible, because if they move at all in sympathy with the woofer excursions it'll mess with how the low end of the mix is represented. How exactly you decide to mount the

PASSIVE RADIATORS AND TRANSMISSION LINES

Not all monitor speakers can be categorized clearly as ported or unported, and Mackie's popular HR-series are a case in point, because although their cabinets are closed, they include a passive radiator (a sort of dummy speaker cone that vibrates in sympathy with the woofer) to achieve ported-style bass enhancement. Another midway design is PMC's transmission-line system, whereby the external port hole feeds a damped internal ducting network designed to reduce the problematic side effects of porting. However, the waterfall plots in Figure 1.4 suggest to me that these strategies are only of limited use in overcoming the resonance issues of ported designs, a suspicion that has been borne out in my own personal experience of several of the different Mackie models.

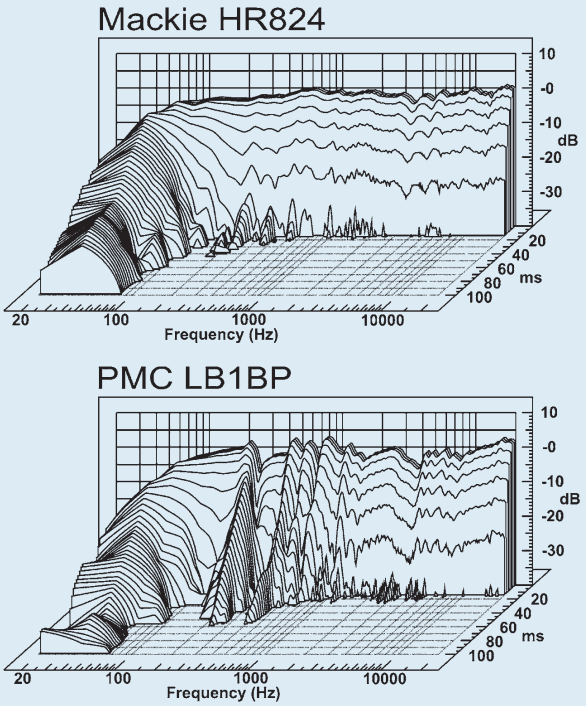

FIGURE 1.4

Waterfall plots for the Mackie HR824 and PMC LB1BP monitors.

boxes will depend on the physical limitations you have to work with in your particular setup, but my recommendation is to use dedicated speaker stands, as these typically give a much better sound than desks and shelves and can be moved around the room more easily than heavy-duty wall brackets. Stands don't need to be exorbitantly expensive either, as long as they are solid enough to keep the speaker still. In fact, you can easily build decent ones yourself if you're handy with your woodwork and use suitably chunky raw materials.

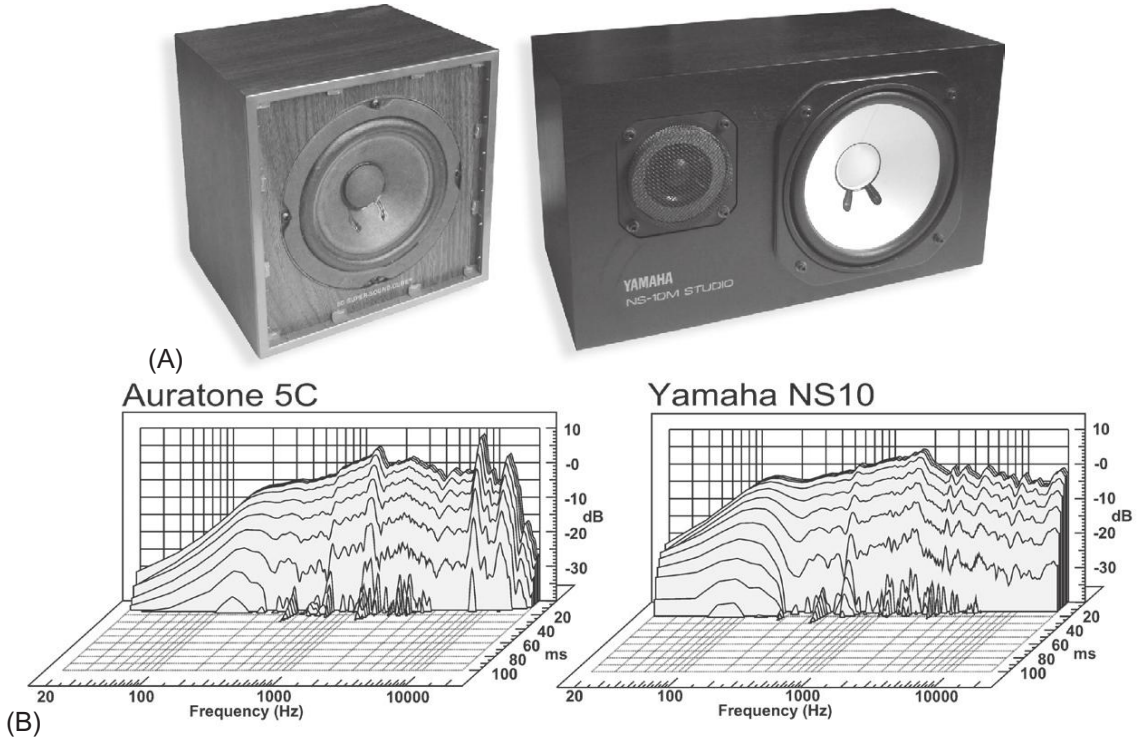

FIGURE 1.5

Two of the most revered mixing speakers are unported designs: the Auratone 5C Super Sound Cube (*left*) and the Yamaha NS10 (*right*). Below them you can see their waterfall plots.

The first thing the mounting hardware has to do is present as much inertia as possible, so the speaker cabinet moves as little as possible in reaction to woofer excursions. If you think the stands/brackets themselves aren't man enough for this job, it's worth trying to add weight to them in the first instance by putting a piece of paving slab underneath each speaker. A bit of rubber matting can also help, by improving the mounting platform's grip on the speaker cabinet. The other main thing your speaker mountings need to do is minimize the transmission of the speaker's physical vibrations into other resonant objects. If your speaker sets off a physical resonance within its stand, for example, it can skew your perception of the mix just as much as any resonances inherent in the speaker design itself. This is one reason why putting speakers on domestic shelves or desks can cause insurmountable monitoring difficulties—it's often surprising how readily these vibrate in sympathy. The LFSineTones audio file is good for revealing resonances, so play it back at a reasonably high volume to see if you hear any furniture humming along. Try also resting a finger on your mounting hardware (stand, bracket, or whatever) and check if you feel any obvious vibrations.

One advantage of dedicated speaker stands is that they often have a hollow frame. It can be filled with sand, which is good for damping resonances and also

increases the stand's inertia. However, other solutions to the resonance problem include foam platforms wedged between the speaker and the mounting surface (such as Auralex's Mo Pads) or little inverted mounting spikes in a similar location (for example Sound Network's China Cones). Primacoustic's Recoil Stabilizer is another popular option, as it incorporates not only a foam base but also a rubberized steel top platform for extra grip and inertia. However, although these kinds of studio widgets can yield noticeable improvements if your monitors are sitting on a normal shelf or desk, my feeling about them is that they're unlikely to justify the extra outlay if you've already sensibly invested £100 ($150) or so in sand-filled speaker stands or similarly sturdy mounting hardware.

1.2 POSITIONING THE SPEAKERS

Whatever you actually sit the speakers on, their exact positioning is also critical to getting good audio reproduction. You should try wherever possible to aim the speakers directly at the listening position. A speaker's frequency response is measured on axis (i.e., from directly in front of it), so if you listen off axis, you won't be hearing what the designer intended you to—high frequencies are more directional than low frequencies, so high-end details in particular tend to suffer. Moving around your listening room should amply demonstrate these effects with any full-bandwidth music mix, but if you want to hear the phenomenon at its starkest, then try listening to a constant full-range test signal such as my PinkNoise file through just one of your speakers. These aren't just miniscule sonic niceties we're talking about. High frequencies are also easily shadowed by physical objects, so make sure you can actually see the drivers you're listening to.

Aiming the speakers isn't just about the horizontal plane either, because vertical alignment is usually even more important, for a couple of reasons. The first is that on most nearfield monitors, the cabinet is profiled around the tweeter to create what's called a waveguide, which is designed to horizontally disperse the jet of high frequencies more widely and thereby increase the size of the optimum listening area (or "sweet spot"). Although waveguides can be quite effective at this, they don't usually do the same job for the vertical high-frequency dispersion and can even make it narrower. But the second reason is that most nearfield monitors have more than one driver in them, with each driver in a different vertical position. A dedicated bit of circuitry or DSP (called a crossover) within the speaker splits the incoming signal's frequency range between the different drivers at factory-specified boundaries (called crossover frequencies). Although ideally the crossover should therefore prevent any overlap between the frequency output of the different drivers, the truth is that there is inevitably a small spectral region around each crossover frequency where two drivers are both contributing significant levels at the same time. If the distance from each driver to the listening position isn't the same, then the signals

FIGURE 1.6
Primacoustic's Recoil Stabilizer cleverly deals with a number of small-studio monitor setup issues simultaneously: the rubberized top mat grips the speaker firmly and couples it to a heavy metal plate, increasing inertia, while the foam base decouples the speaker from physical resonances in the underlying surface.

from the different drivers will arrive at the listening position at different times (or "out of phase" in geek-speak), and this gives rise to a potentially serious frequency-cancellation effect called comb filtering.

PHASE AND COMB FILTERING

I've touched on the ideas of phase and comb filtering in the main text, but because they have so many ramifications when mixing, it's worth looking at the subject more closely. The best way to start thinking about phase is first to consider a sine-wave signal, the simple audio waveform from which all complex musical sounds can theoretically be built. A sine wave generates only a single audio frequency, according to how many times its waveform shape repeats in a second. For example, a 1kHz sine wave repeats its waveform 1000 times per second, with each waveform repetition lasting 1ms. Imagine that you have two mixer channels, each fed from the same sine-wave source at the same frequency. The peaks and troughs of the two waveforms will be exactly in line, and mixing them together will simply produce the same sine wave, only louder. In this situation we talk about the two sine waves being "in phase" with each other.

If you gradually delay the audio going through the second channel, however, the peaks and troughs of the two sine waves shift out of alignment. Because of the unique properties of sine waves, the combination of the two channels will now still produce a sine wave of the same frequency, but its level will be lower than if the two channels were in phase, and we say that "partial phase cancellation" has occurred. When the second channel is delayed such that its peaks coincide exactly with the first channel's troughs (and vice versa), the two waveforms will combine to produce silence. At this point we say that the waveforms are completely "out of phase" with each other and that "total phase cancellation" has occurred.

When total phase cancellation occurs, you sometimes hear engineers say that the signals are "180 degrees out of phase." This is a phrase that's not always used correctly, and it can therefore be a bit confusing. To describe the phase relationship between two identical waveforms, mathematicians often quantify the offset between them in degrees, where 360 degrees equals the duration of each waveform repetition. Therefore, a zero-degree phase relationship between two sine waves makes them perfectly in phase, whereas a 180-degree phase relationship puts them perfectly out of phase, resulting in total phase cancellation. All the other possible phase relationships put the waveforms partially out of phase with each other, resulting in partial phase cancellation. What's confusing about the "180 degrees out of phase" term is that it is sometimes used to refer to a situation where the second channel's waveform has been flipped upside down, so that the peaks become troughs and vice versa—a process more unambiguously referred to as polarity reversal. This scenario also results in silence at the combined output, hence the common confusion in terminology, but it's very important to realize that the total phase cancellation here is brought about by inverting one of the waveforms, not by delaying it.

Now let's scale things back up to deal with real-world sounds, made up as they are of heaps of different sine waves at different frequencies, each one fading in and out as pitches and timbres change. If we feed, say, a drum loop to our two mixer channels, instead of a single sine wave, any delay in the second channel will have a dramatic effect on the tonality of the combined signal, rather than just altering its level. This is because for a given delay, the phase relationships between sine waves on the first channel and those on the second channel depend on the frequency of each individual

(*Continued*)

sine wave. So, for example, a 0.5ms delay in the second channel will put any 1kHz sine-wave components (the waveforms of which repeat every 1ms) completely out of phase with those on the first channel, resulting in total phase cancellation. On the other hand, any 2kHz sine-wave components (the waveforms of which repeat every 0.5ms) will remain perfectly in phase. As the frequency of the sine-wave components increases from 1kHz to 2kHz, the total phase cancellation becomes only partial, and the level increases toward the perfect phase alignment at 2kHz.

Of course, above 2kHz the sine-wave components begin partially phase canceling again, and if you're quick with your mental arithmetic you'll have spotted that total phase cancellation will also occur at 3kHz, 5kHz, 7kHz, and so on up the frequency spectrum, whereas at 4kHz, 6kHz, 8Hz, and so on the sine-wave components will be exactly in phase. This produces a characteristic series of regularly spaced peaks and troughs in the combined frequency response of our drum loop—an effect called comb filtering. A delay of just 0.000025s (a 40th of a millisecond) between the two channels will cause total phase cancellation at 20kHz, but you'll also hear partial phase cancellation at frequencies below this. As the delay increases, the comb filter response marches further down the frequency spectrum, trailing its pattern of peaks and troughs behind it, which themselves get closer and closer together. However, when the delay times reach beyond about 25ms or so (depending on the sound in question), our ears start to discern the higher frequencies of the delayed signal as distinct echoes, rather than as a timbral change, and as the delay time increases, phase cancellation is restricted to progressively lower frequencies.

Although it should now be clear that the tonal effects of comb filtering can be disastrous if two identical signals are combined with a delay between them, most real-world comb filtering at mixdown is actually much less severe, either because the out-of-phase signals aren't completely identical, or because they're at very different levels, or both.

Although manufacturers typically do their best to keep crossover regions pretty narrow to minimize the effect of comb filtering, most affordable nearfield monitors have only two drivers, which means that any comb filtering between the woofer and the tweeter happens in the worst possible place from a mixing standpoint: right in the center of the frequency spectrum, where our hearing is most sensitive. If you want to get a handle on the extent of the damage here, try this experiment. Play the PinkNoise file (*www.cambridge-mt.com/ms-ch1.htm*) through a single nearfield speaker with vertically spaced drivers, and listen to it first of all directly on axis. Now drift alternately about six inches to each side while keeping your vertical position constant. You'll hear a small change in tone on most speakers because of the high-frequency directionality I mentioned earlier. Once you're used to that change, drift up and down by about six inches instead, and the tonal change will likely be much more noticeable. Although the effects of comb filtering between your speaker drivers won't

Most affordable nearfield monitors have only two drivers, which means any comb filtering between the woofer and the tweeter happens in the worst possible place from a mixing standpoint: right in the center of the frequency spectrum, where our hearing is most sensitive.

usually be as obviously apparent in their own right when you're listening to a real-world mix, that doesn't mean they aren't there, and the ripples they put into the frequency response treacherously undermine your ability to judge both the tone and level balance of critical sounds in the midrange—things like lead vocals, snare drums, and guitars.

Stereo Monitoring

A small studio's nearfield monitors will usually provide the most reliable source of information about a mix's stereo image, but for them to do this effectively they have to be set up so that the distance between the speakers equals the distance from each of the speakers to the listening position. This is because the human brain has evolved a powerful instinct that causes us to perceive the source of an incoming sound as being located on the same side as the ear that sound reaches first. (Presumably natural selection favored those cavemen who could work out where the saber-toothed tiger growls were coming from.) This instinct means that you don't have to move your head very far out of the sweet spot before the majority of the stereo image folds down into the closer speaker.

Pulling the speakers too far apart is a frequent mistake too, as it destabilizes the center of the stereo image, making balance judgments for the important sounds that are typically placed there tricky. If anything, it's better to err on the side of placing the speakers too close together, because the narrowed stereo picture this produces is a lot easier to work with at mixdown than unstable central imaging. You'll also get the best stereo image if you try to position the whole monitoring setup such that the room is fairly symmetrical around the listener's line of sight, in order to retain image balance despite the impact of any sonic reflections from room boundaries and furniture.

SPEAKERS ON THEIR SIDES?

There is a persistent myth among small-studio owners that putting speakers on their sides is the more "pro" method. True, a quick web surf will furnish you with countless pictures of *Battlestar Galactica*–style control rooms in which nearfield speakers are visible perched sideways on top of the console's meterbridge, but that setup has little to do with monitoring fidelity and everything to do with getting the nearfields out of the way of the big main monitors and maximizing the view through the control-room window during recording dates.

If you separate your speaker's drivers horizontally by placing the cabinets on their sides, then you have to be much more careful to keep your horizontal listening position absolutely consistent (both side-to-side and front-to-back) if you're going to avoid being stung by crossover comb-filtering effects. With vertical speakers, on the other hand, you'll only get these crossover problems if you move you head vertically, which gives you greater freedom of movement while you're working.

Add to this that the waveguides in most nearfield monitors are designed to broaden the high-frequency sweet spot when the speaker is vertical. If you flip that waveguide on its side it narrows the sweet spot instead, as well as bouncing more high-frequency energy off nearby desk and ceiling surfaces—something that's not sensible from an acoustics point of view.

One final pitfall to avoid is having your speakers out of polarity with each other. Normally any sound at the center of the stereo image should be coming out of both speakers at an equal level and with all the speaker drivers moving back and forth in sync. However, it's surprisingly common for less experienced studio users to mistakenly wire up their rig such that one driver is pushing toward the listener while the other is pulling away, such that the left-channel waveform is effectively inverted compared to the right-channel waveform, a situation usually referred to as having your speakers "out of phase," or to be more accurate "out of polarity." (To clarify: two audio waveforms are out of phase when there's a time delay between them; they're out of polarity when one of the waveforms is inverted, such that peaks become troughs and vice versa.) If you've got this problem, you'll get a very odd stereo listening experience, which feels a bit like having your brains sucked out your ear, and it also makes both stereo positioning and level balance difficult to judge.

To check for this effect, listen to the StereoTest audio file (*www.cambridge-mt.com/ms-ch1.htm*) through your system. It contains a repeating pattern of four noise-burst test signals: the first only in the left channel, the second only in the right channel, the third in both channels, and fourth in both channels, but with the right channel out of polarity with the left channel. The former pair of noise bursts will confirm that your speakers are indeed operating in stereo and are connected the right way round; the latter pair of noise bursts should make it fairly obvious if your speakers are out of polarity with each other; the third burst should be much more clearly central in the stereo image than the fourth if all is well. If you do discover that your speakers are out of polarity, then the finger of blame will almost certainly point to some facet of your audio wiring after the outputs of your studio's playback system, the most common culprit being that the positive and negative terminals of passive speakers haven't been connected correctly with the respective terminals on their amplifier.

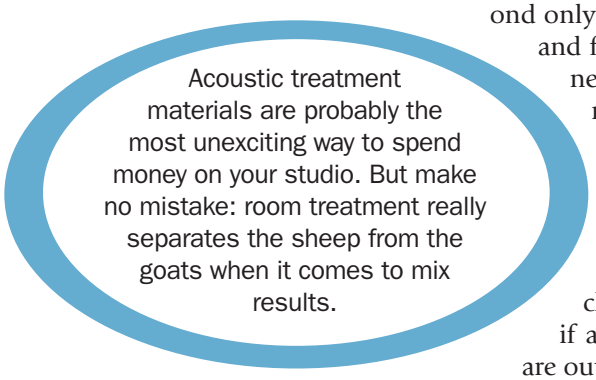

1.3 DEALING WITH ACOUSTIC REFLECTIONS

Room acoustics is one area of small-studio design that's woefully neglected in the main. It's not that budding engineers don't pay lip service to the worthiness of acoustic design in principle, it's just that acoustic treatment materials are probably the most unexciting way to spend money on your studio. They make no sound. They have no twinkly lights. They can give you splinters. But make no mistake: room treatment really separates the sheep from the goats when it comes to mix results. Spike Stent minces no words: "You can have the best equipment in the world in your control room, but if the room sounds

like shit, you're onto a hiding to nothing."[4] The room acoustics are at least as important to the sound of your monitoring as your speaker hardware is. Permit me to shout that again for you: THE ROOM IS AT LEAST AS IMPORTANT AS THE SPEAKERS! So that means engineers looking for a fast track to commercial-sounding mixes should plough as much money into dealing with their room acoustics as they do into buying their monitors.

Now I'm the first person to acknowledge that few people can afford to engage specialist acousticians to build their studio from the ground up. Indeed, my experience is that small setups almost universally find themselves squished unceremoniously into a shed, basement, spare bedroom, attic, or office unit, where the user faces a battle even to adjust the interior decor, let alone the construction and placement of the walls. However, even if you're relegated to the cupboard under the stairs, it's still vital that you make what acoustics improvements you can. The good news here is that you can work wonders with off-the-shelf products and a bit of do-it-yourself (DIY), even on a limited budget and in rental accommodation where your construction options are limited. When you consider that a typical untreated domestic acoustic environment will, in my experience, render roughly two thirds of the money you spent on your speakers wasted, there's simply no excuse for inaction if you're serious about your craft.

The first main set of acoustics problems that you'll need to lock horns with when working on nearfield monitors is that sound doesn't just squirt directly out of the front of each speaker and then stop when it hits your ears. It also sprays around in all other directions to a greater or lesser extent, bouncing off the walls, ceiling, and floor, as well as any other reasonably solid object in the room. This can create a situation where a sonic reflection arrives at the mixing position shortly after the sound that traveled there directly—in other words, you end up listening to two versions of roughly the same sound slightly out of phase. As I mentioned in Section 1.2, this kind of combination of out-of-phase signals can cause nasty comb-filtering ripples in the perceived frequency response—and not just in your speakers' relatively restricted crossover regions, but across the entire frequency spectrum.

The most troublesome reflectors you'll encounter will be the furniture you're working on, whether that's a large mixing console or the bedside table propping up your laptop, as well as any room boundaries within about three meters of your listening position. (To be honest, sonic reflections coming from farther than three meters away aren't a primary concern in domestic environments, because their greater time delays and diminished levels are unlikely to cause major comb-filtering problems.) In professional setups the room surfaces around the monitoring position are often cleverly angled to fire all the main sound reflections from the speakers away from the listening position, but that kind of approach is rarely viable in the smaller rectangular rooms of more modest studios. A more practical alternative in most cases is to place absorptive

USING A SUBWOOFER

There are certain styles of music for which the very lowest frequencies are extremely important, so if that's the market you're working in then you'll need a speaker system that can actually reproduce them. According to Trina Shoemaker, "Forty Hertz exists now.... We have real low end in today's recordings, so you have to work with it."[5] One common way of extending the low-frequency reach of budget nearfield systems is to supplement a smaller pair of stereo "satellite" speakers with a single additional subwoofer to create a so-called 2.1 system. Although this would seem to compromise the stereo presentation of the mix, we actually hear stereo primarily from higher frequencies, so that's not a problem in practice. In fact, you can usually position subwoofers well off-center without unbalancing the stereo picture. An advantage of the 2.1 approach is that you don't get phase-cancellation issues between drivers at the low end, and you also get a certain amount of flexibility to position the subwoofer independently of the nearfield satellites to reduce room-mode problems (a serious acoustics issue we'll discuss in detail in Section 1.4). However, although many manufacturers of 2.1 systems suggest that you can place the subwoofer pretty much wherever you like, I'd personally advise restricting yourself to placements where the subwoofer and satellites are all an equal distance from your ears; otherwise you'll compromise the relative timing of low-frequency transients. Also, as with any speaker, beware of porting side effects on cheaper products.

acoustic treatment at all the main reflection points in order to reduce the level of the reflected sound, and hence the severity of the comb filtering.

Acoustic Foam in Moderation

A cost-effective absorber here is open-cell acoustic foam of the type manufactured by companies like Auralex, and a meter-square area of such treatment covering each of the reflection points can make a big difference to the clarity of the speaker tone. In a typical small setup, that means you'd put patches of treatment on each wall, on the ceiling, on the wall behind the monitors, and on the wall behind the listener if that's within range. Some further foam on your work surface can be a bonus too, but clearly your options may be limited here! In general, the thicker your absorptive material, the lower down the frequency spectrum it's going to be able to absorb, so don't just use the cheapest foam you can get, because it'll probably only be 5cm thick or less; a 10cm thickness will do a much more respectable job. If you have to cut treatment costs, try using a pattern of thinner and thicker foam tiles, rather than going for thinner foam across the board.

You can easily glue acoustic foam directly to walls or ceilings with an appropriate spray adhesive, but if you do that you'll have the devil's own job getting it off again if you ever want to transplant your studio setup at a later date. It's therefore a much better idea to glue the foam to some kind of lightweight plyboard instead, so you can then hang the foam like a picture in whichever room

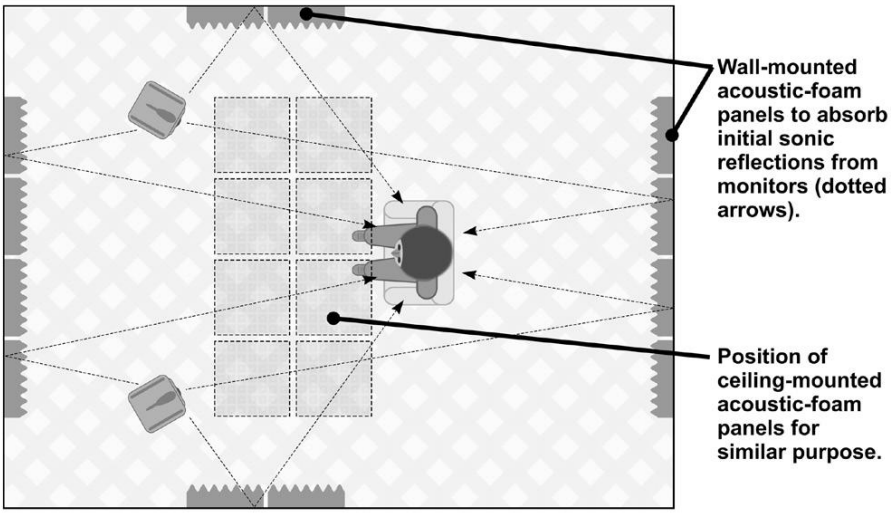

FIGURE 1.7
If sound from your speakers reflects off control-room surfaces back toward the monitoring sweet spot, it can cause all sorts of frequency-response problems for the listener. This diagram shows the wall reflection points for a typical small rectangular control room and how you can use acoustic foam panels to reduce the problem. Don't forget that the ceiling reflects sound too!

you happen to want to work in—a great advantage if you have to get results in untreated college facilities or on other people's home rigs. You could also try angling the foamed boards to bounce some of the reflected frequencies away from the sweet spot, just as in purpose-designed rooms.

Above all, however, resist the temptation to cover the entire room in foam—I can't tell you how often college studios in particular succumb to this seductive error. It is quite simply a recipe for disaster, because it hoovers up the top half of the room's reverberation, making for an extremely unnatural working environment. Although it makes good sense to damp down strong early reflections that can potentially comb filter your frequency response, you also want your monitoring environment to bear at least a passing resemblance to real-world living rooms and offices. If you plaster your whole studio in foam, you'll basically be mixing for people in padded cells—perhaps not the most lucrative demographic to set your sights on! What's more, the economics of covering such a large surface forces most foam fanatics to compromise on treatment thickness, which just sucks the life out of the upper frequencies, giving a superficial impression of acoustic control while the rest of the spectrum merrily runs riot.

If you plaster your whole studio in foam, you'll basically be mixing for people in padded cells—perhaps not the most lucrative demographic to set your sights on!

"I've heard a lot of home recordings," says Keith Olsen, "and there's a lot of what they think is deadening, because it takes out all the top end. That might save the neighbors from calling the cops, but the bottom end and midrange is still real ambient; it bounces around the room, and you get phase destruction."[6]

Because acoustic absorption like this is best used in moderation, it shouldn't break the bank—a couple of hundred pounds is about as much as you need, or indeed want, to spend for foam in a small studio. And if you're sensibly dividing your budget equally between the speakers and the acoustics, then even entry-level speakers should justify this kind of outlay. Still, if for whatever reason you can't afford proper acoustic treatment, you'll find that even soft furnishings such as thick curtains, blankets, and duvets can be of some benefit in damping those reflections if rigged appropriately. One more tip in this instance, though, would be to try to leave a few inches of air gap between the curtains/blankets and your room boundaries, as that has a broadly similar effect to increasing the thickness of the treatment. (You can pull this stunt with acoustic foam too by sticking small foam spacer blocks behind the main foam panels, as long as you're fairly open to the idea of covering yourself in glue and swearing a lot during the process.)

Boundary Effects

There's one further reflection issue to be wary of: a constructive low-frequency interference commonly referred to as the boundary effect. As you move a speaker closer to a room boundary, it reduces the delay between the direct and reflected sounds arriving at the listening position, making them less and less out of phase. This means that the comb filtering reduces and you start to get just a simple level reinforcement as the reflected sound adds power to the direct sound. However, for two reasons this reinforcement occurs primarily at low frequencies: first, their longer wavelengths shift less out of phase as a result of a given delay, and second, low frequencies are better transmitted by the speaker off-axis anyway. So if you place your speakers right up next to a wall, you'll get up to 3dB of bass boost, and this can rise to 6dB if you tuck them into a room corner where the effects of two boundaries gang up.

One solution is to EQ the speaker's output to compensate for the low-end tip-up. Indeed, a lot of active monitors aimed at compact setups have a little low-cut switch round the back for just this purpose. However, although this is one of the only situations where EQ can usefully bail out your acoustics, I'd still advise against placing your speakers right up against a wall if you can help it, because even with a sensible thickness of acoustic foam on that surface, there is still likely to be enough reflected sound arriving at the listening position to give significant comb-filtering problems in the midrange. Furthermore, if you're using monitors with ports at the rear of the cabinet, the proximity of the boundary is likely to increase turbulence as air zips in and out of the port opening, leading to extra noise and low-frequency distortion.

WHAT ABOUT DIFFUSION?

Another way to reduce the comb-filtering impact of early reflections is to use acoustic diffusers on reflective surfaces to scatter the reflection in lots of different directions. The downside of using diffusers in the small studio, though, is that they're usually more expensive than acoustic foam for a given area of coverage. However, that doesn't mean that diffusion has no part to play in project studios, because it turns out that things like CD racks and bookshelves can work quite well, as long as they're stocked fairly irregularly. (If ever there were a good excuse for having a messy studio, then that's it!) Eric Schilling is a big fan of using a bookshelf like this: "It has mass, and each book has a different depth and size. The concept is brilliant in its simplicity."[7] So do try to position shelves like these usefully if you can—the wall behind the monitoring position is a particularly good bet, because that otherwise takes a fairly large area of acoustic foam to treat effectively, which risks overdeadening the room's high-frequency reverberation.

1.4 TACKLING ROOM RESONANCES

Although acoustic reflection problems can make mincemeat of monitoring accuracy, the remedies I've suggested are cost-effective, fairly simple to implement, and effective enough that comb filtering shouldn't stand in the way of you achieving commercial-level mixes. It's hardly surprising, then, that the more switched-on small-studio owners have often already implemented something along these lines. However, there is another equally important aspect of room acoustics, which is more difficult to tackle and so is often simply ignored by budget-conscious operators: room resonance.

Understanding the Problem

To understand how room resonances work, it helps to bring to mind how a guitar string resonates. At its lowest resonant frequency (called the first mode), the string is stationary at both ends and moves most at its middle point—or to use the technical terms, there are nodes at the ends of the string and an antinode in the middle. However, the string also has a second resonant mode at twice this frequency, giving a vibration with three nodes, such that the string appears to be vibrating in two equal-length sections. Tripling the first mode's frequency gives you a third mode with four nodes, quadrupling it gives you a fourth mode with five nodes, and so on up the spectrum.

The reason it's useful to keep this image in mind is that the body of air between any parallel room boundaries has a similar series of resonant modes at frequencies dictated by the distance between the surfaces. A quick-and-dirty way to work out the resonant frequency of the first room mode between a given pair of parallel boundaries is by dividing 172 by the distance in meters between them. Subsequent room modes will then be at multiples of that frequency, just as in our guitar-string example. So if the ceiling of your studio is 2.42m above the floor, then you'd expect the first room mode in that dimension to be at around 71Hz, the second at 142Hz, the third at 213Hz, and so forth.

A single room mode can easily push its resonant frequency 20dB out of kilter, so only a flying pig is likely to find a good listening position when several room modes are active at the same time.

Each room mode will generate its own regularly spaced series of nodes and antinodes between the room boundaries, and if this means that there's a node in your monitoring sweet spot, you'll hear a drastic frequency-response dip at that room mode's resonant frequency, whereas if there's an antinode at the listening position, you'll hear a significant level boost at that frequency instead. Because each pair of parallel room surfaces will contribute its own independent series or room modes, and most rectangular domestic rooms offer three pairs of parallel surfaces, small studios typically find themselves liberally peppered with nodes and antinodes at different frequencies.

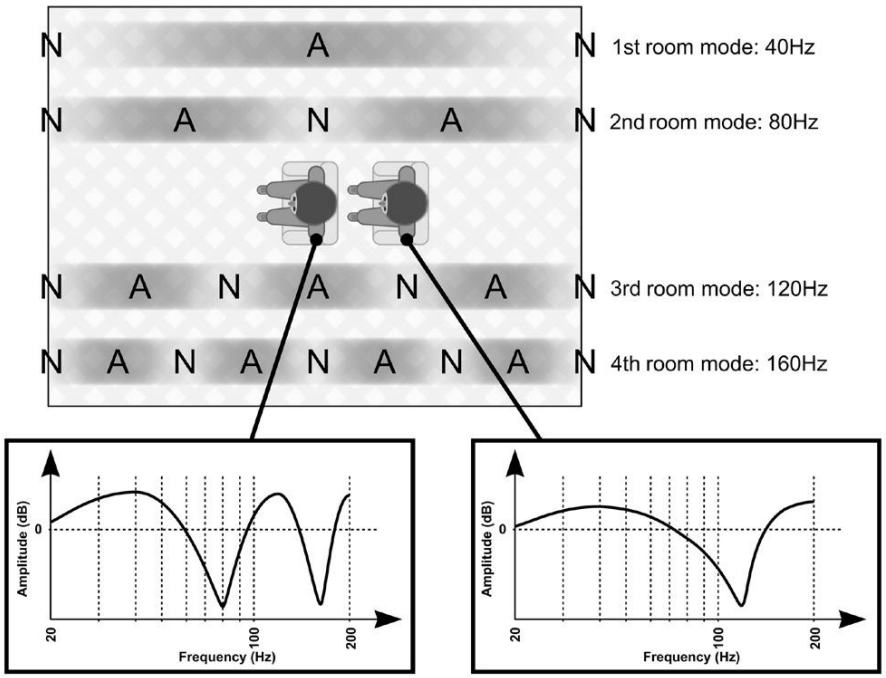

FIGURE 1.8

This diagram demonstrates what room resonances can do to your monitoring system's apparent frequency response. This picture shows the first four front-back room modes for a room measuring around 4.3m long. These occur at 40Hz, 80Hz, 120Hz, and 160Hz. The nodes and antinodes are marked with "N" and "A," respectively, and although they are shown one above the other here for the sake of clarity, it's important to realize that they are actually overlaid on one another, all occurring simultaneously across the whole width of the room. The two frequency-response plots demonstrate the effects of these modes on the monitoring system's apparent frequency response at two different listening positions only about 75cm apart.

So what does this mean in practice? Well, the first thing to say is that even a single room mode can easily push its resonant frequency 20dB out of kilter, so only a flying pig is likely to find a listening position that gives a faithful spectral balance when several room modes are active at the same time. Plus, if you move around the room at all while listening, the monitoring system's apparent frequency response will writhe around like a python doing the lambada—as I've tried to illustrate with the frequency plots in Figure 1.8. To be fair, room modes tend to affect primarily the lower half of the audio spectrum, by virtue of the fact that higher-frequency resonances are much more easily damped by normal room decor, but the remaining sub-1kHz disaster area is more than enough to scupper your hopes of making objective decisions about a mix.

Every room is different, though, so try this experiment to get a realistic idea of what the room modes are doing to your own monitoring. Play back the LFSineTones audio file through your system again and listen carefully from the sweet spot, comparing the relative levels of the pure sine-wave tones as they march in semitone steps up the bottom three octaves of the audio spectrum. If your studio is anything like most small, untreated control rooms, you'll probably find that some tones almost disappear, whereas others practically go into orbit! Table 1.1 shows roughly which frequencies occur at which times in the file, so grab a pencil and make a note of which are the most wayward tones while you're listening at the sweet spot. Now move your listening position a couple of feet away and try that little experiment again—it'll be a whole different story, and you'll probably find that some of the tones that were anemic before are now overprominent, and vice versa.

Now it would be quite reasonable to say that sine-wave tones aren't much like a real musical signal, so it's also worthwhile to focus on what your room is doing to the evenness of bass lines on commercial tracks that you know to have been very well-produced in this department. If you want a suggestion here, then try the song "All Four Seasons," produced and engineered by Hugh Padgham for Sting's album *Mercury Falling*. The bass part on this track is wide ranging, but it is also extremely consistent, so all the notes should sound fairly even on any mix-ready monitoring system. If they don't, then you've seriously got to ask yourself how you're planning to judge bass balances in your own mixing work.

Table 1.1 LFSineTones Audio File Map

Track Time	Frequency	Pitch
0:00	24Hz	F
0:01	25Hz	F#
0:02	26Hz	G
0:03	27Hz	G#
0:04	28Hz	A
0:05	29Hz	A#
0:06	31Hz	B
0:07	33Hz	C
0:09	35Hz	C#
0:10	37Hz	D
0:11	39Hz	D#
0:12	41Hz	E
0:13	44Hz	F
0:14	47Hz	F#
0:15	49Hz	G
0:16	52Hz	G#
0:17	55Hz	A
0:18	59Hz	A#
0:19	62Hz	B
0:20	65Hz	C
0:22	69Hz	C#
0:23	73Hz	D
0:24	77Hz	D#
0:25	82Hz	E
0:26	87Hz	F
0:27	92Hz	F#
0:28	98Hz	G
0:29	105Hz	G#
0:30	111Hz	A
0:31	117Hz	A#
0:32	123Hz	B
0:33	131Hz	C
0:35	139Hz	C#
0:36	147Hz	D
0:37	156Hz	D#
0:38	165Hz	E
0:39	175Hz	F
0:40	185Hz	F#
0:41	196Hz	G
0:42	208Hz	G#
0:43	220Hz	A
0:44	233Hz	A#
0:45	247Hz	B
0:46	262Hz	C

Some Practical Remedies

The professional solution to low-end resonance problems is building rooms with nonparallel walls, and if you've got builders in the family then by all means follow that route. However, for the rest of us the choice of which room we choose to mix in can also make a big difference. For a start, it's sensible to avoid small rooms where possible, because their resonances trespass further up the frequency spectrum than those of larger spaces. Also try to find a room where the dimensions aren't matched too closely, otherwise multiple similar room modes will gang up on the same frequency ranges, and that's just asking for trouble. Some home studios come spectacularly unstuck here, because they've been banished to a small 2.5m-wide cubic spare room where the powerful combination of room modes in all three dimensions conspires to make a complete dog's dinner not just of the low end, but also of the midrange.

Whatever room you're in, you can reduce the impact of the room modes if you avoid setting up your listening position (or indeed your monitors) exactly halfway between any of the room boundaries, where you're likely to get the worst pileup of nodes and antinodes. However, as I've already mentioned, off-center setups bring with them the danger of a lopsided stereo balance because of unmatched room reflections, so it's advisable not to push the sweet spot too far left or right of center. That said, a stereo imbalance causes far fewer mixing problems in my experience than room-mode issues, so I'd personally prioritize room-mode treatment if push comes to shove.

Rooms with lightweight walls can work to your advantage by allowing more low frequencies to escape, rather than resonating within—assuming that your neighbors don't take a dim view of this sound leakage! By the same token, concrete-walled basement rooms should be approached with caution, because low frequencies can have real trouble escaping that kind of environment and you'll have your work cut out trying to tame the resulting room modes. However, most small studios simply have to take whatever room happens to be free, so in any case you need to know how to make the best of things using acoustic treatment.

Mineral-Fiber Bass Traps

The best all-purpose tactic is to damp down the room modes as much as you can using low-frequency absorbers, often called bass traps. The downside here, though, is that bass traps need to be dense and bulky to do their job properly. As Eric Schilling notes, foam simply isn't up to the task: "Most people think that treating a room simply means going to a music store and buying foam. But if it's essentially a square room, it doesn't matter if you have some foam in the corner and a few pieces on the wall—you still won't be able to hear bass to save your life!"[8] The most commonly used alternative is large slabs of high-density mineral fiber, which offer much better low-frequency absorption. Placing the panels close to a given room boundary provides broadband absorption of all the associated dimension's room modes, and (much as with

CAN EQUALIZATION CORRECT THE EFFECTS OF ROOM MODES?

Because room modes cause low-end frequency-response changes, it's tempting to think that EQ might be able to offer a solution to this kind of resonance problem—not least because a number of manufacturers now offer EQ-based "room correction" software ostensibly for that purpose. The idea with these algorithms is that they measure the monitoring system's frequency response using a special test signal and a calibrated microphone, and then calculate an EQ curve to attempt to compensate for the nonlinearities they detect. However, for anyone serious about mixing on a budget, I think such systems are a red herring, for two different reasons. First, staying exactly in the sweet spot all the time isn't actually conducive to getting work done in a real studio, and if you wander out of the sweet spot, the frequency response will quickly go out of alignment again, as demonstrated in Figure 1.8. Even if you can somehow clamp yourself into the sweet spot, anyone working in the room with you will hear something completely different. Plus, if you're stuck with having to use your control room for recording as well, the fluctuating frequency response as you move round the room will make finding decent miking positions a tedious guessing game.

More crucially, though, room resonances don't just affect frequency response; they also cause specific frequencies to ring on in time, with all the detrimental consequences we've already explored in relation to monitor porting back in Section 1.1. Equalizers are themselves resonant too, which only compounds the issue. So can equalization usefully combat room modes? For my money, no.

foam) you get an increase in the low-frequency absorption if you leave an air gap behind the slab—a foot or so if possible. Because mineral-fiber bass traps are typically more expensive than foam panels, one common trick is to place them across wall-to-wall and wall-to-ceiling corners, where they can treat modes in two dimensions simultaneously, and in small rooms this has the added advantage of intruding less on your workspace.

Normally not all dimensions of a room are equally troublesome, so it's not a bad idea to try to work out which are the biggest culprits. To do this, first listen carefully to the LFSineTones audio file (*www.cambridge-mt.com/ms.ch1.htm*) while referring to Table 1.1, taking care to identify the most problematic frequencies—anywhere you can hear a big peak or trough in the apparent levels of the tones. Then divide 172 by each of these frequencies in turn to give a list of measurements in meters, and look with suspicion on any room dimension that works out as a simple multiple of any of those measurements. Once you know which dimensions are causing the biggest problems, you can concentrate your acoustic treatment resources more effectively.

Companies such as Primacoustic, Ready Acoustics, and Real Traps all offer ready-made bass traps incorporating high-density mineral fiber, but the small studios I've been into usually need around a dozen 10cm-thick slabs to make their monitoring workable, and with off-the-shelf products that can seem

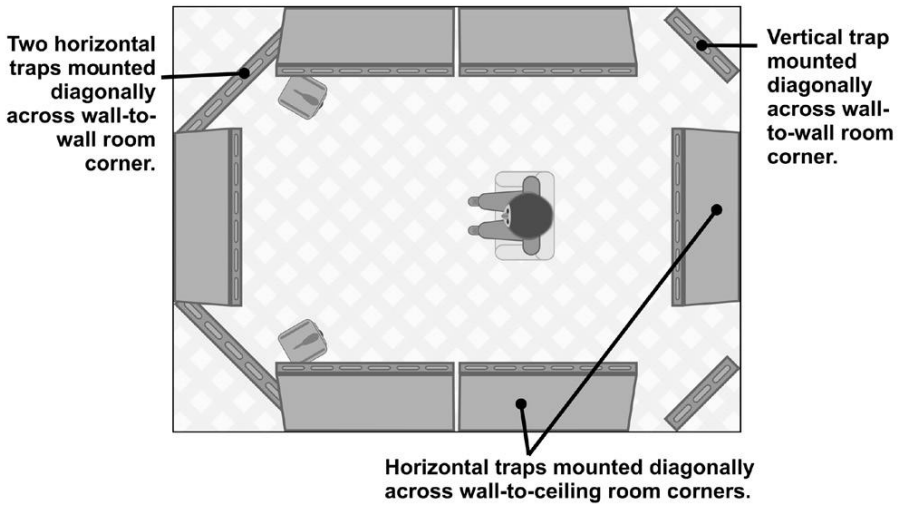

FIGURE 1.9
This diagram shows how you might sensibly distribute a dozen 10cm × 60cm × 120cm mineral-fiber acoustic panels within a modestly sized control room to tame typical room-mode problems.

rather an unpalatable battering of one's credit card. That's understandable, though, argues Mick Glossop: "Controlling very low frequencies costs more money, because it's big movements of air, and you need mass and big trapping systems to do that."[9]

However, the good news is that a bit of DIY can save you serious money, because it's possible to build reasonable bass trapping on your own at a fraction of the cost. What you need to find are 10cm-thick mineral-fiber panels with a density of around 50 to 100kg/m^3. They're used for insulation as much as for acoustics, so ask at an insulation company if you can't locate a dedicated supplier of acoustics materials in your area. A word of caution, though: mineral fiber is a skin and breathing irritant, so be sure to wear appropriate protective mask, goggles, and gloves when handling it. By the same token, you'll want to cover each panel in some kind of acoustically neutral fabric to stop the mineral fibers from shedding all over the place—the kind of material commonly used to cover office dividers is quite a good bet, especially because you can usually get it in a wide range of colors. Ready Acoustics also do little ready-made zipped fabric sacks, called Ready Bags, which are another fairly cost-effective alternative if you're less than handy with a needle and thread, and they're neat looking to boot.

A bit of DIY can save you serious money, because it's possible to build reasonable bass trapping on your own at a fraction of the cost of off-the-shelf products.

Because high-density mineral-fiber panels are about as stiff as a typical foam mattress, you can just use picture hooks and string to hold them in place in your studio if you're happy for

them to sag a little over time under their own weight. If you want a tidier look, though, then there's nothing stopping you from building a simple wooden frame around them (either inside or outside the fabric), and covering it in chicken wire to hold the slab in place more assertively. Alternatively, if you're as much of a danger to yourself with a saw in your hand as I am, then check out Ready Acoustic's Chameleon Acoustic Frames, which fit around standard 10cm × 60cm × 120cm slabs to create a slick final product.

TOO MUCH BASS TRAPPING?

Although I've heaped scorn on the idea of covering every available control-room surface in acoustic foam, there's little benefit in similar restraint when it comes to bass trapping in small studio rooms—the more the merrier. But hang on a minute! Didn't I say earlier that too much acoustic foam could make a control room less like a real-world listening environment? Shouldn't that also apply to bass trapping? Well, yes it does. The difference, however, is that leaving your control room's modes to their own devices (as in most real-world playback environments) will actually make your low-end monitoring less representative, because any other room will have a different set of modes at a different set of frequencies. Better to tame your control-room modes as well as you can so that you've a clearer idea of the actual low-frequency balance, irrespective of any resonances of the end user's listening system.

One practical problem you may encounter when installing a lot of mineral-fiber bass trapping, though, is that it will also absorb high frequencies as well and can result in too dead a sound in that region of the spectrum—much the same problem as you get when too much acoustic foam has been installed. In these cases you may actually need to reintroduce some high-frequency reflection by fixing small areas of hard surface to the faces of some of the bass traps. Small picture frames, old CD-ROMs, or just random bits of hardboard are all possibilities here, although you should avoid covering more than about a third of the trap's surface area in this way or else you'll start to interfere unduly with its low-end absorption.

Limp-Mass Bass Traps

Now mineral-fiber slabs can do a lot of good, but even quite thick ones become less effective as you head down below 100Hz, so in a solidly built room with powerful sub-100Hz resonances, there's a limit to how much they can help. For example, there's one 6m × 4m × 2.5m basement room I've used that resonated powerfully at 30Hz and 60Hz along its long front-back dimension, despite the presence of two dozen 10cm × 60cm × 120cm mineral-fiber bass traps. In such cases, there's one other type of treatment that can be worth a shot: a limp-mass trap. This involves loosely hanging a large, heavy, and impermeable sheet a little distance from one room boundary, where it can damp low-frequency air movements. Some kind of rubberized barrier matting with a density of around 5kg/m^2 is a common choice for this task, and it helps if it has a reinforced backing so that you can if necessary hang it by its edge at ceiling level without it tearing

or deforming under its own considerable weight. Failing that, old carpets can deliver something like the same result if mounted in a similar way and may be a cheaper option if you can get hold of them as scrap.

Lower frequencies need larger areas of treatment, so if you feel the need for limp-mass trapping for a troublesome room mode in one particular dimension, then you should think in terms of trying to treat pretty much the whole of one of the relevant room boundaries, even though this will inevitably reduce the area you can use for your studio gear. Because the size of air gap behind the trapping adjusts its absorptive properties, it's a good idea to mount the matting on some kind of movable wooden frame if possible, so that you can use trial and error to strike the best balance between resonance reduction and loss of workspace. This kind of trapping is a lot less predictable than simple mineral-fiber absorbers, because it is itself to some extent resonant, so be prepared to spend a bit of time refining the setup to get the best out of it. Some variations worth considering are adjusting the fixing points for the matting, as well as hanging drapes, thin boards, or mineral-fiber panels in parallel. It's not an exact science, but if you're faced with heinous low-end resonance problems and a meager budget, then it can nonetheless be a viable bacon-saver. In the specific basement room I mentioned earlier, putting in free-hanging sheets of barrier matting across most of the width of the room and about a meter away from the rear wall was able to bring the worst low-end problems under control, and the loss of that workspace was a small price to pay for usable monitoring.

FLUTTER ECHO

Although the main resonance problems in studio rooms are at low frequencies, you can also get higher-frequency resonances too, often referred to as flutter echoes. The simplest test for these is clapping your hands in the sweet spot and listening for any hint of unnatural metallic-sounding buzz as the sound dies away in the room. If you're suspicious, then another patch of acoustic foam on one or both of the offending parallel surfaces should put an end to it with minimal fuss because high frequencies are so easily absorbed.

1.5 WHEN IS MY MONITORING GOOD ENOUGH?

There's little doubt that the standard of the nearfield monitoring system is one of the most significant quality bottlenecks in the small-studio production process. "The biggest mistakes that happen in project studios," says Al Schmitt, "are the result of the monitoring systems."[10] For this reason, improving your monitoring is perhaps the smartest way to spend money, and no investment in this department is likely to be wasted. Nevertheless, studio owners only ever have so much money to go around, so they are understandably keen to spend only as much as they really need in order to reliably achieve commercial-level mix results. In this spirit, let me suggest the kind of system that, in my

experience, can deliver that level of performance while remaining on the right side of the point of diminishing returns.

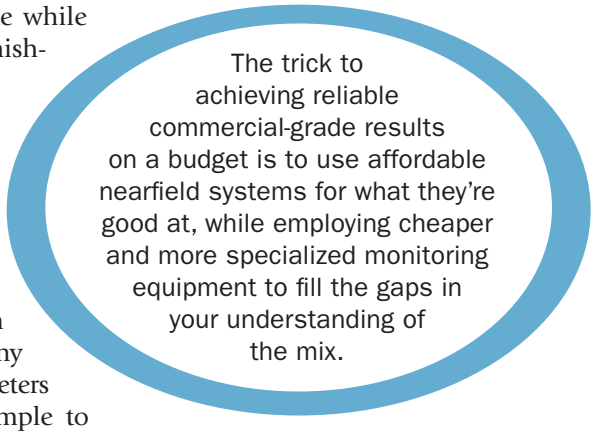

As I've already mentioned, ported nearfields only tend to become really useful for mixing above the £1000 ($1500) mark, whereas there are cheaper unported systems that hit this usefulness level, as long as you don't need to stun small prey with your playback volume. A pair of solid speaker stands filled with sand will do such a set of speakers justice without the need for any extra mounting gizmos, and four or five square meters of 10cm thick acoustic foam will usually be ample to bring early reflections and flutter echoes under control. Unless you have solid concrete room boundaries or your studio is stuck in a basement, a dozen 10cm × 60cm × 120cm mineral-fiber bass traps should be enough to give you usably even bass response as well. Basically, with that lot in the bag your nearfield monitoring shouldn't be any excuse for turning in duff mix work.

Now before I'm lynched by a posse of high-end speaker designers, acousticians, and mastering engineers, let me be clear that the setup I'm recommending here will still be far from perfect. But perfection is not the point of it, because the trick to achieving reliable commercial-grade results on a budget is to use affordable nearfield systems for what they're good at, while employing cheaper and more specialized monitoring equipment to fill the gaps in your understanding of the mix. Indeed, if the nearfield rig I've suggested is still way beyond your current means and you have no alternative but to work with hopelessly compromised nearfield monitoring, the truth of the matter is that you can still achieve surprisingly good mix results by relying more heavily on such supplementary monitoring tools. So what are these additional bits of gear? All will be explained in the next chapter, but before we get to that, let's quickly recap my main recommendations so far.

CUT TO THE CHASE

- A nearfield monitoring system is a good choice for small-studio mixing. Spend as much as you can afford on the speakers, because quality costs, and if your budget is tight then be wary of ported designs. When choosing a system, favor studio monitors over hi-fi speakers, active models over passive, and accuracy over volume.
- Whatever speakers you use, mount them securely on solid, nonresonant surfaces, preferably away from room boundaries. If the speakers have more than one driver, then the cabinets should be oriented so that the drivers are equal distances from the listener and angled toward the listening position. In nearly all cases it's better for multidriver speakers to be vertically rather than horizontally aligned. For stereo listening, there should be the

same distance between the listener and each of the speakers as between the speakers themselves, but if anything err on the side of a narrower speaker spacing. Make sure you check your speakers for polarity.

- Give some thought to the room you use for your studio, and ensure that you spend at least as much money on acoustic treatment as you do on monitors. Strategic use of acoustic foam can effectively target early-reflection and flutter-echo problems, but be careful not to overdo it. High-density mineral-fiber slabs can provide a fairly foolproof remedy for low-frequency room resonances, but if these don't prove effective enough in your room, then try supplementing those with additional limp-mass trapping. Don't waste your time trying to correct acoustics problems with equalization, because the benefits of this are minimal in practice.

Assignment

- Invest as much money as you can in your nearfield speaker system, spending roughly the same amount on acoustic treatment as on the speakers themselves.
- Make the best of whatever system you can afford (or have access to) by making sure that the speakers are solidly mounted and sensibly positioned and that the room is appropriately treated.

www.cambridge-mt.com/ms-ch1.htm

CHAPTER 2

Supplementary Monitoring

As I touched on at the end of Chapter 1, aiming for perfect nearfield monitoring in the small studio is a fool's errand. With the best will in the world, you'll always have problems, not least because it's neither practical nor comfortable to confine yourself to a tiny listening sweet spot for hours on end in order to maintain the best stereo imaging, minimize comb-filtering problems between the speaker drivers, and stay clear of room-mode danger zones. However, even if you've got some medieval torture instrument clamping your head into the optimal location, you'll still face a problem inherent in all stereo systems—namely, that your speaker cabinets are positioned at the left and right extremes of the stereo field, whereas many of the most important parts of your mix will be center stage, where there is no actual hardware. In other words, your lead vocals, kick, snare, and bass will all typically appear to be hovering in thin air in front of you, an illusion called a "phantom image," which is achieved by feeding equal levels of each of these tracks to both speakers at once.

The difficulty with phantom images as far as real-world mixing is concerned is that they invariably feel less stable than the "real" images at the extremes. Part of the reason for this is that normal manufacturing tolerances prevent any pair of remotely affordable speakers from being exactly matched as far as frequency response is concerned, which smears the frequency spectrum of centrally panned sounds haphazardly across the stereo image to a small but significant extent. The effects of typical small-room resonance modes and asymmetrical early reflections only make things worse, and then there's the fact that small movements of the listener's head will affect the sound of phantom images much more than that of real images because of comb filtering between the two speakers.

So what's the extent of the damage in your room? Fire up the StereoTest audio file again (*www.cambridge-mt.com/ms-ch2.htm*), and listen to how the "real" left/right images of the first two noise bursts are narrower and more precise than the phantom central image of the third burst. (You can ignore the fourth burst in the pattern for this test.) You'll hear this effect to some extent even on high-spec monitoring systems, but in budget setups it's like night and day.

Mixing Secrets for the Small Studio.

A real-world musical example that demonstrates this same thing quite well is the track "American Life" from Madonna's album of the same name. Here the verse lead vocal is very narrow and panned dead center, whereas the synth that appears 45 seconds in and the acoustic guitar that appears 15 seconds later are panned to opposite stereo extremes. Every small-studio system I've ever clapped ears on renders the synth and guitar images much more stable than the vocal's.

The big drawback with carefully configured studio nearfields is that they don't actually sound anything like the majority of real-world playback systems.

The final big drawback with carefully configured studio nearfields is that they don't actually sound anything like the majority of real-world playback systems, so you can't expect them to tell you how your mix will fare once it's out in the big wide world. Yes, you should get a fairly good idea what multiplex popcorn munchers and chin-stroking hi-fi tweeds will hear, but you'll be pretty much in the dark regarding the low-fidelity, low-bandwidth playback of most mainstream listening devices.

So the most important thing to remember with your nearfields is to rely on them for what they're best at:

- For investigating what's going on across the whole frequency spectrum, especially at the frequency extremes
- For judging the impact of mix processing on the sheer quality of your sounds
- For evaluating and adjusting the stereo image
- For understanding how your mix will sound on more high-fidelity listening systems
- For impressing the pants off the artist, manager, or A&R rep

For a number of other monitoring requirements, however, there are more modest systems that can actually outperform even a decent small-studio nearfield rig, simply by virtue of being better adapted to specific duties. This chapter introduces some of these additional systems, specifically the ones you can't afford to do without if you're seriously planning on competing with the big guns. Let's start by looking at the most powerful of these, as epitomized by one of the most famous mixing speakers in studio history, the Auratone 5C Super Sound Cube.

2.1 THE GHOST OF THE AURATONE 5C SUPER SOUND CUBE

Despite the Auratone 5C's unassuming appearance and subjectively unappealing sound, it's been pivotal to the mixes of numerous hit records, including Michael Jackson's *Thriller*, the biggest-selling album of all time. "I love Auratones!" enthuses Jackson's longtime producer and engineer Bruce Swedien.

"You know what Quincy [Jones] calls them? The Truth Speakers! There's no hype with an Auratone.... Probably 80 percent of the mix is done on Auratones, and then I'll have a final listen or two on the big speakers."[1] Auratones also mixed the soundtrack to *Saturday Night Fever* (for many years the best-selling film soundtrack album), and you can still rely on most large studios having at least one of them on hand for those freelance engineers who appreciate what they have to offer—no matter how star-spangled the main and nearfield systems. These are all persuasive reasons to bring the magic of the Auratone to bear in your own mixing work.

FIGURE 2.1
The best-selling album ever, Michael Jackson's *Thriller*, was primarily mixed on Auratones.

But that's where there's a big old fly in the ointment: Auratones aren't manufactured any more, and much of the remaining secondhand stock has now succumbed to the frailties of old age. So what can you do? Well, as long as you understand the nature of this speaker's extraordinary mixing abilities, and the design features responsible for them, then you can use this knowledge to track down a modern replacement that offers similar benefits.

Midrange Focus

First things first: the Auratone is portless, and we've already discussed at length the implications of that in Chapter 1. The next crucial

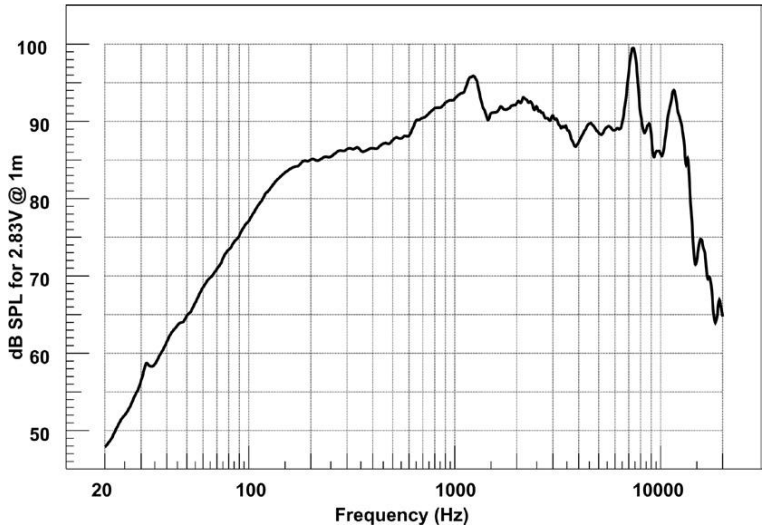

FIGURE 2.2
The frequency response of the Auratone 5C Super Sound Cube, which clearly shows its midrange bias.

attribute is the single small driver, and for several reasons. First, its restricted ability to reproduce both high and low frequencies focuses your attention on the midrange of your mix, which is the frequency region our ears are most sensitive to and which tends to carry the most important musical information. "The real perspective lives in that range," says Jack Joseph Puig. "It doesn't live in the highs, it doesn't live in the lows. That's what really speaks to the heart."[2] Bob Katz concurs: "The midrange is the key. If you lose the midrange, you lose it all."[3]

The midrange is also crucial because it's the frequency region of your mix that is most likely to survive the journey to your target listener's ears. "The midrange is... most common to every system in the world," says Puig. "Some system might have five woofers, one might not have any, one might have ten tweeters, but all of them have midrange, and so *that* has got to be right."[4] For one thing, you need a big listening system to produce much in the way of low frequencies, so small real-world playback devices (iPods, clock radios, phones, and a lot of TVs and hi-fis too) won't give you much output in the sub-200Hz region. A huge number of mass-market playback systems also share the Auratone's one-cone design, which inherently detracts from high-frequency efficiency. Just think of all the small boomboxes, TVs, and multimedia speakers of this type that litter most offices and households. Now add in the arguably more widespread multispeaker announcement and piped-music systems that assail you while you're blowing froth off your skinny latte, stepping aboard a 747 to Malaga, or pressing the button for the 17th floor—more than one driver per speaker in large audio distribution networks like these would seriously bump up the installation cost, so it's pretty rare.

But it's not just an end-user's hardware that compromises the frequency extremes; it's the way it's used, too. If you recall that high frequencies are very directional, easily shadowed, and swiftly absorbed by normal soft furnishings, it's little surprise that extended high frequencies rarely trouble the public, simply because few people actually listen to speakers on axis. In fact, it's more than likely that their hi-fi speakers are concealed behind the sofa or their portable radio is tucked away on that top shelf at the other side of the office. At the other end of the spectrum, the high levels of low-frequency background noise in cars, planes, public spaces, and some workplaces will obscure most of that portion of a mix. (Yes, I know that some people pimp their rides with weapons-grade subwoofers, but those guys are probably mostly deaf into the bargain.) It should be clear by now why the Auratone's design is so adept at demonstrating which bits of your mix will reach the greatest number of listeners. In a nutshell, if your mix doesn't stand up against the competition on an Auratone, then you've not done your job properly!

Resistance to Acoustics and Comb-filtering Problems

There are other practical advantages to the Auratone's single-driver construction that are especially relevant to small studios. The first is that the restricted

low-frequency output barely tickles the room modes in most cases, with the result that low-end resonances don't interfere nearly as much with your mix decisions, especially when it comes to balancing bass instruments. The speaker isn't actually very loud overall either, and that reduces the impact of early reflections from the room as well.

Another real benefit of the single-cone design is that no crossover electronics are required to split the signal between different speaker drivers. Any processing between your mixer's output and your speakers has the potential to introduce side effects which can mislead you while mixing, and in the case of crossovers the worst gremlins tend to manifest themselves around the crossover frequency. Most affordable nearfield monitors have their crossovers in the 2 to 3kHz midrange region, so the razor-sharp midrange focus of a crossoverless Auratone provides a valuable alternate perspective.

By the same token, Auratones won't punish you with interdriver comb filtering if you slouch a bit in your studio chair. What's more, if you restrict yourself to using just a single Auratone, rather than a pair, you'll also get no comb filtering from side-to-side movements, so you're free to stroll around the control room with comparative impunity—you'll get some high-end loss off-axis, but that actually does very little to undermine the Auratone's essential usefulness.

Balance Judgments and Mono Compatibility

Of course, if you want to listen to a stereo mix from just the one Auratone, then it forces you to listen back in mono. Rather than being a disadvantage, though, this only adds to the appeal. For a start, single-speaker mono suffers

FIGURE 2.3
A large number of real-world speakers have only one driver, including shopping-center piped-music systems, small radios/TVs, and mobile phones.

none of the phantom-image instability that plagues even decent nearfield stereo setups, simply because it doesn't create any phantom images at all—the physical location of the speaker is unequivocally the source for everything you hear. Central sounds therefore feel just as solid as those at the stereo extremes, so it becomes possible to judge exact level balances in your mix with pinpoint accuracy, regardless of stereo positioning.

However, you need to be aware that summing the two channels of a stereo mix to mono shifts the overall balance of the sounds. Partly this is an inevitable side effect of the way that stereo speakers create their illusory panorama—two sounds that appear equally loud in stereo will actually appear to be at different volumes in mono if they're positioned at different distances from the center of the stereo image. A good rule of thumb here is that if you listen to a stereo mix in mono, then the levels of central sounds will increase by roughly 3dB relative to sounds located at the edges of the picture. To hear this effect in action on your own system, check out the MonoBalanceShift audio file (*www.cambridge-mt.com/ms-ch2.htm*). It's made up of a pattern of four short brass-band chords that are differentiated only by their position in the stereo image: chord 1 is central; chord 2 is at the left-hand extreme; chord 3 is back in the center; chord 4 is at the right-hand extreme; and then the four-chord pattern starts again from chord 1. The levels of the chords are such that their loudness should appear to remain fairly constant as they trot backward and forward across the stereo field—assuming that you've set up your monitoring system roughly along the lines suggested in Chapter 1. However, if you now switch to monitoring in mono, you'll hear a clear louder-softer alternation in the chord levels.

FIGURE 2.4
Stereo recordings made with spaced-pair microphone techniques are a common reason for mono-compatibility problems.

The fact that the mono balance shift draws more of your attention to what's going on in the center of the stereo image can be a blessing in disguise, though—after all, the most meaningful sounds in the mix usually reside there. Mono listening also forces you to work harder to achieve clarity for each of the sounds in your mix, because you can't make things more audible just by shifting them to less cluttered areas of the stereo field. As Geoff Emerick puts it, "Mixing in stereo [is] the easy way out."[5]

Mono listening forces you to work harder to achieve clarity for each of the sounds in your mix, because you can't make things more audible just by shifting them to less cluttered areas of the stereo field.

However, you can definitely get some unpleasant surprises when switching to mono if your stereo mix is sneakily harboring any phase or polarity problems between its left and right channels. A common source of these kinds of nasties is stereo recordings that have used spaced microphone configurations—drums, room ambience, vocal/instrumental ensembles, and piano are often captured in this way, for instance. If any part of the sound source isn't the same distance from both mics, its sound will arrive at one mic earlier than at the other. In stereo this time-delay between the two channels actually reinforces the stereo positioning of the instrument for the listener, but in mono the same time delay will cause comb filtering, the frequency-cancellation effect that we encountered several times in Chapter 1. To hear an example of this in practice, listen to the DrumOverheadsStereoPhase audio file (*www.cambridge-mt.com/ms-ch2.htm*), which is a short section of a raw spaced-pair drum-overhead recording, presented alternately in stereo and summed mono versions. As the file switches to mono there is a clear loss of high end on the cymbals and hi-hat in particular, making the sound as a whole rather boxy. In addition to real stereo recordings, simulated stereo reverb and echo effects can also suffer from comb filtering in mono, because they combine delayed versions of the same basic sound by their very nature.

If any signal appears in both stereo channels, but with inverted waveform polarity (i.e., with the waveform shape inverted) in one of them, then that can also cause a serious mismatch between the balance of the stereo and mono renditions of your mix. Fortunately, most normal recorded sounds don't have much in the way of out-of-polarity stereo components—the most common exception being stereo recordings where a rogue cable or signal processor has inadvertently flipped the polarity of just one of the mics. However, a lot of synthesizer patches deliberately do incorporate one-sided polarity inversion, because that gives them a subjectively impressive artificial wideness in the stereo image. (Some dedicated stereo-widening mix effects work along similar lines too.) The drawback of this widening trick, though, is that out-of-polarity stereo components simply cancel themselves out when a stereo mix is summed to mono, which can lead to unwelcome surprises for any important synth hooks in your arrangement. Listen to the SynthPadStereoPhase audio example (*www.cambridge-mt.com/ms-ch2.htm*) to get an idea of how dramatic a balance shift this kind of cancellation can cause. The

file presents the same small section of a stereo synth sound in both mono and stereo versions, and the subjective level difference between the two versions is enormous. Now imagine that was your track's lead synth hook, and smell the fear!

Is Mono Still Relevant?

But surely now that recordings, downloads, and broadcasts are all routinely in stereo, mono is only really relevant to a few grunting Neanderthals, right? Wrong. Although the majority of music is indeed now produced, sold, and transmitted in stereo, you might as well kiss goodbye to most of the stereo information in your mix once it's been through the end user's reception or playback equipment. Even for the most earnestly bearded hi-fi enthusiast, stereo imaging will still crumple into the nearest speaker if he shuffles out of the sweet spot and, frankly, I can't actually recall ever having seen a home hi-fi system with anything worth calling a sweet spot anyway. Hands up—how many of you have acoustic treatment in your living room? Exactly. And how many of you have had the speakers banished by your flatmates to that corner under the side table, behind the plant? You're lucky to hear the lyrics from most hi-fis, let alone any stereo.

FIGURE 2.5
Two common real-world speaker setups that prevent any meaningful stereo imaging reaching the listener: hi-fi speakers placed artistically in a CD rack (*left*) and speakers hidden away in high corners of shops and restaurants (*right*).

Less fancy domestic playback systems fare even less well. The left speaker of your teenager's hi-fi mini-system is far too busy holding the wardrobe door open to worry about stereo imaging, even if it weren't accidentally wired up with inverted polarity. And it stands to reason that your web PC's speakers have to huddle together at one side of the monitor to leave room on the desk for the phone. Furthermore, watching any TV from the sweet spot between its two built-in speakers is a surefire route to eyestrain. Maybe you can save the proper stereo experience for the kitchen boombox, building a custom shoulder harness to keep you securely between the built-in minispeakers while you peel those spuds and heat up the fryer? Perhaps not. (Are you even sure both speakers are working? My kitchen radio recently presented me with a version of "Yesterday" minus the strings.) To be fair, earbuds are still something of a safe haven for stereo, but only if you resist the urge to share one of them with the person sitting next to you and you're not one of the millions of listeners worldwide who, like Brian Wilson, suffer from one-sided hearing impairment.

Move outside the home and the situation is no less dire. Consider the lengths to which shop and restaurant owners seem to go to avoid any risk of stereo imaging

interrupting your retail experience. "No matter what anyone says," says Allen Sides, "if you're in a bar you're going to hear one speaker."[6] Car stereo? I only know of one extortionately priced sports car that actually puts the driver an equal distance from the two speakers, so anyone else wanting a true stereo experience in the car will have to run the gauntlet of a potentially eye-watering encounter with the gear shift. But at least in these cases you have different speakers playing different things. Many other systems don't even make a token stab at stereo! Things like telephone hold systems, band/club PAs, and shopping-center announcement systems all typically sum their inputs to mono before feeding any speakers. FM radio receivers also often automatically sum the audio to mono in order to improve reception in the presence of a weak transmitted signal. With all this in mind, checking the mono-compatibility of your mix remains just as indispensable as it was 50 years ago, unless you're unfazed by the thought of your audience being treated to a hideously imbalanced and comb-filtered version of your latest track during their weekly supermarket sweep.

An Auratone won't impress any clients with its pathetic honky sound, but it's nonetheless a brilliant specialist tool that ruthlessly spotlights midrange balance and mono-compatibility issues in your mix.

Modern Auratone Substitutes

Allow me to recap here. If you want to reap the greatest benefits of the Auratone in your mixing environment, then you need to monitor in mono from one small, unported, single-driver speaker. A lot of people still harbor the misguided view that the Auratone's primary feature was that it sounded horrible (hence its common "Horrortone" nickname); in other words that it simply allowed you to anticipate the effects of worst-case scenario playback. However, this myth doesn't really stand up to much scrutiny when you consider that the Auratone's portless cabinet is unusual in small speakers these days, and that it also delivers unusually low distortion. Clearly, a single Auratone won't impress any of your clients with its pathetic honky sound, but it's nonetheless a brilliant specialist tool that ruthlessly spotlights midrange balance and mono-compatibility issues in your mix. As with your main nearfield system, you just need to learn to use an Auratone for what it excels at while relying on your other monitoring equipment to make up for its weaknesses.

FIGURE 2.6
A modern homage to the Auratone 5C: Avantone's Mix Cube.

Finding a modern speaker to fulfill the Auratone role isn't that easy, because most designs have a port, two drivers, or a single driver that is too small. Probably the clearest Auratone substitute available as a current product is Avantone's Mix Cube. This is a great little speaker, which sees daily use in my own mix system and which in some respects actually improves on the

original Auratone's specs. Although at £205 ($235) per active speaker the Mix Cube would top my list in terms of price/performance ratio, there are nonetheless a few other affordable options I've discovered: Pyramid's Triple P, Fostex's 6301, and Canford Audio's Diecast Speaker. All but the Triple P are available with amplification built in.

An ancillary issue is how to feed the speaker in mono. Thankfully, increasing numbers of fairly affordable mixers and monitor controllers now incorporate the required mono and speaker-selection buttons, but if your own system is lacking these then there are lots of workarounds. For example, software mixing systems will usually let you route a mono version of the mix to a spare audio interface output, so you can feed your Auratone from that output and then switch between speakers from within the software. If you're listening via a hardware mixer, then route the mix through a spare pair of channels (or a dedicated stereo channel) and feed the Auratone using a mono prefader auxiliary send. (Just make sure that these channels don't also route back into your main mix buss or you'll get a potentially speaker-busting feedback howlaround!) The mixer's internal send routing should take care of the mono summing for

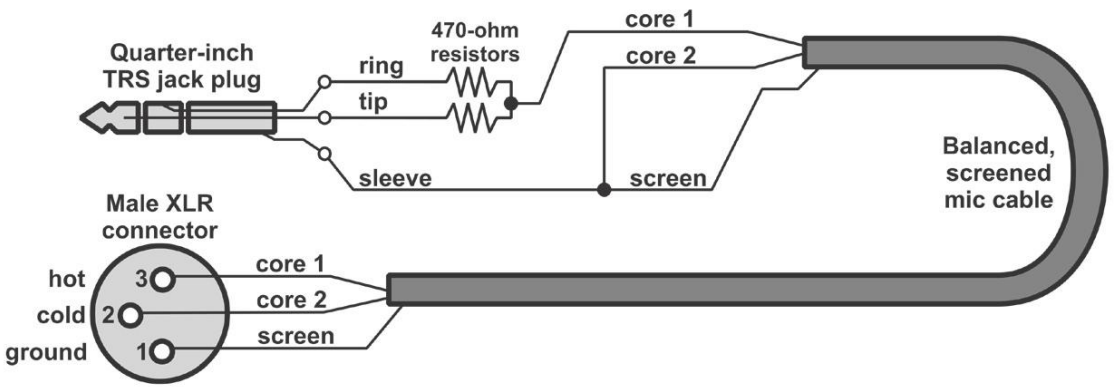

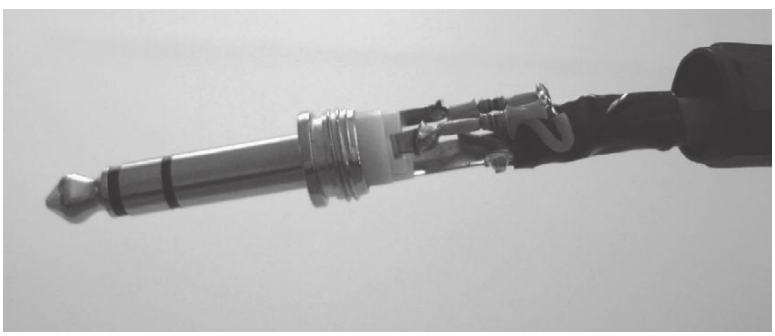

FIGURE 2.7

Schematic for creating a lead that will feed mono to an Auratone-substitute from a stereo headphone socket (*above*) and an example of the finished product (*below*).

you, and when you want to hear the Auratone you just mute the mix buss and crank up the relevant auxiliary send controls. Neither of these workarounds will help, though, if you're working on a small system with a lone pair of stereo output sockets. Fortunately, in such cases there's often still an additional stereo headphone output that can be pressed into service if you're handy with a soldering iron. Figure 2.7 shows how you can put together a cable that sums the two channels of a headphone output to mono for this purpose.

2.2 HEADPHONES

Since the Sony Walkman first arrived on the scene back in 1979, headphone listening has become so widespread that you'd have to be pretty cavalier to sign off a mix without checking how it translates for such an enormous group of consumers. Therefore, the second important supplementary monitoring system I'd suggest using as a matter of course is headphones. Clearly there are limits to how much low-frequency output headphones can deliver, but your Auratone should already have alerted you to any problems that might arise from bass-light playback, so that shouldn't be news to you. What's more important about headphones is that almost all of them transmit each side of your mix exclusively to one ear, whereas with speakers the two stereo channels are always heard to some extent by both ears. For this reason the stereo image is much wider on headphones (the image encompassing an angle of 180 degrees, where speakers only cover around 60 degrees), and any sound at the stereo extremes feels disconnected from the rest of the mix as a result. Whether the image stretching that headphone listeners experience helps or hinders your particular mix will depend on your own preferences and the expectations of the musical genre you're working in, but if you don't check your mix on headphones, then none of your decisions will be properly informed.

But the usefulness of headphones for monitoring in the small studio isn't limited to telling you how a sizeable chunk of the public will perceive your mix. Headphones also serve an additional practical purpose, because their sound isn't affected nearly as much by your monitoring environment. In the first instance, background noise from computer fans, nearby traffic, or next door's line-dancing class are reduced in level relative to the mix signal when using headphones, which means that fewer subtle technical problems will slip through the net. Things like brief audio dropouts, bad edits, and momentary fader-automation slipups stand out much more clearly on headphones, and it's also easier to detect the onset of digital clipping on stereo files—the left and right channels will usually clip at different times, so the sprinkles of digital distortion appear characteristically right at the edges of the 180-degree stereo field. The lack of room reflections when monitoring on headphones is another advantage, because it serves as something of a safety net in the event that comb filtering or room-resonance problems are seriously jeopardizing your balance judgments. If you've done your Chapter 1 homework, then you shouldn't need to rely on this precaution heavily, but if you're working on an unfamiliar system or have only limited control over your

WHAT, NO NS10S?!

I mentioned in Chapter 1 that Yamaha's NS10s are far and away the most celebrated mixing speakers in the professional music world, and two of the main reasons for this also apply to the Auratone: the sealed-box cabinet (with its smooth low-frequency roll-off, excellent transient response, and low distortion) and the midrange-heavy frequency balance (which reflects the tonal deficiencies of many small real-world playback systems). However, if you want the maximum bang for your buck in the small studio, I wouldn't personally recommend investing in a pair of these classic speakers for yourself. First of all, NS10s won't really give you much more useful information than the two monitoring systems I've so far discussed: a reasonable full-range nearfield system should easily match the capabilities of NS10s when it comes to judging stereo imaging and tonal quality, whereas an Auratone will deliver much the same closed-box accuracy and midrange emphasis. The second reason for considering the NS10 a nonessential purchase is that it doesn't obviate the need for either a decent nearfield system or an Auratone. It can't deliver the frequency extremes in the way reasonable nearfield systems now can (which is presumably why Chris Lord-Alge[7] uses his NS10s with an additional subwoofer), and its mix-balancing abilities can't match those of a single Auratone because of the same phasing complications that afflict all multidriver stereo speaker systems. Now I'm not saying that the NS10s don't deserve their kingpin status; it's just that they aren't as well-suited to the small studio as they are to the large commercial setups where they beautifully bridge the gap between main wall-fitted monster monitors and any Auratone/grotbox in use.

own monitoring environment, then a pair of familiar-sounding headphones can be a godsend.

If your budget's tight, then you'll be pleased to know that the audio quality required of a set of headphones in this context isn't tremendous, and £50 ($75) should give you all the fidelity you need, but do still try to get studio (rather than hi-fi) models if you can. In fact, the kinds of headphones that most recording musicians use while overdubbing are usually more than up to the mark. That said, there are actually strong reasons why most small-studio owners may find it worthwhile to spend more money here if they can. The first is that top-of-the-range studio monitoring headphones are capable of doing many of the jobs you'd expect of speakers, which is great for situations where you only have limited access to decent nearfields. You might be one of 20 students jostling for time in your college's single studio; perhaps you can't afford to set up your own decent nearfield system yet and can only check things periodically on a mate's rig; maybe you've been threatened with slow lingering death if you wake that baby one more time. Regardless of the reasons, the ability to get a measure of real mixing done on headphones can dramatically increase the amount of time you have available for improving your production's sonics.

Clearly, the peculiarities of the stereo picture will present some additional panning and balance difficulties whenever you work primarily on headphones, plus even the best headphones won't give you a proper impression of your low end.

On the flipside, though, the absence of room acoustics problems means that you may find that audio-quality and tone decisions across the rest of the spectrum actually turn out to be more dependable. In my experience, the problems when working with excellent headphones are by no means insurmountable as long as you still have at least some access to an Auratone and a pair full-range nearfields—whoever's they happen to be. As a matter of fact, given that my personal favorite top-end headphones (Beyerdynamic's DT880 Pro, Sennheiser's HD650, and Sony's MDR7509 HD) all retail for under £350 ($500) I almost always recommend one of those sets as a first monitoring system for those starting out. My honest opinion is that unless you've set aside at least £1500 ($2000) for nearfield speakers and acoustic treatment, then you'll get more reliable mixes from a pair of top-drawer headphones—particularly if you're able to hijack someone else's nearfield system briefly from time to time to check the low end and stereo image.

So why bother with nearfields at all if headphones can provide much of the same information? As in a lot of studio situations, a big reason for the outlay is speed. It's quicker to do the bulk of your mixing work on a listening system that gives you the whole frequency response in one shot. Although it's perfectly feasible to produce commercial-quality mixes without nearfields of your own, it involves a good deal of extra legwork to sort out the bass end in particular when headphones are the main workhorse, as we'll discuss further in Chapter 3. General mix balancing on headphones also typically takes more time, because

FIGURE 2.8
Top-of-the-range headphones, such as the Beyerdynamic DT880 Pro (*left*) and Sony MDR7509HD (*right*) shown here, are a good investment for most small-studio operators, especially when budget is limited.

the wide stereo image and overexposed mix details can be misleading—common pitfalls are balancing both lead parts and delay/reverb effects too low. To keep yourself on the straight and narrow you'll have to rely more heavily on an Auratone substitute while processing your tracks, and you'll also need to compare your work much more frequently with competing commercial productions. All of this only really makes sense in the long term if you've got a whole lot more time than money. Or to put it another way, a cheap razor blade will cut a blade of grass beautifully, but I still don't begrudge the extra expense of a lawnmower.

2.3 GROTBOXES

I explained in Section 2.1 that although the Auratone has a reputation for sounding nasty, it's not actually a good reference point when it comes to lo-fi playback, so to round out your supplementary monitoring options it makes sense to have at least one set of "grotboxes," unashamedly cheap mass-market speakers that let you hear how your mix holds up in the face of serious playback abuse. It's not difficult to find such speakers—the little ported active ones that you can feed from a PC's headphone output are a good bet, for example, or perhaps the speakers of a bargain-basement boombox. Don't be afraid to drive them fairly loud either, as this'll only emphasize the poorly controlled distortion and resonance artifacts, and it's the impact of these that you want to evaluate.

One advantage of using a pair of separate PC-style speakers is that you can put them right alongside each other without much in the way of a gap between them. This is how a lot of people end up using such speakers anyway, and it's also roughly how the speakers are spaced in small boomboxes. The effect of placing the speakers like this is to remove any useful stereo information while at the same time avoiding centrally placed sounds from increasing in level as you'd expect when summing a stereo mix signal to mono electrically. This is about the worst scenario you can get as far as the clarity of central signals is concerned, so it's the real acid test of lead-vocal level and intelligibility. If you can hear your melody and lyrics clearly under these circumstances, then that's about the best guarantee you have that they'll make it through to the widest possible audience. I also find it handy to put grotbox speakers outside your studio room if possible and to listen to them off-axis from a good distance away. Again, this is a typical mass-market listening experience, and forces you to work even harder to keep the mix clear when broad-based commercial appeal is the aim.

Despite what you might think, grotboxes can frequently be found in the most high-spec of commercial mix rooms. Take Bob Clearmountain's pair of compact Apple computer speakers, for example: "Those are actually my favorites! I have them placed right next to each other at the side of the room on top of a rack. There's just no hype with the little ones—they're so close together that everything is almost mono, and I get a really clear perspective of how the overall thing

sounds.... They're also good for setting vocal levels."[8] Michael Brauer uses something pretty similar too: "Most of my mixing is done listening to my little Sony boombox. It sits behind me... about four-and-a-half feet up from my ears and six feet back, and I listen at a medium low level. Because the speakers of the boombox are relatively close to each other, I essentially listen in mono. The Sony is like a magnifying glass, it tells me whether my mix sucks or not."[9]

Finally, while some engineers like to refer to several different grotboxes during the mix process (boomboxes, TVs, car stereos, you name it), my view is that if you've already used your nearfields, Auratone substitute, and headphones sensibly, then there's little extra information to be gleaned from this, and the extra time and money could be better used in other ways.

FIGURE 2.9
A pair of cheap speakers placed side by side like this is the acid test for lead-vocal levels and intelligibility in your mix.

CUT TO THE CHASE

- A reasonable small-studio nearfield monitoring system can give a good overview of the full frequency spectrum and the stereo field, and it is well-suited to assessing the effects of mix processing on the quality and tone of your mix. It should also give you a good idea of what hi-fi listeners and cinema audiences will hear, and is most likely to impress your clients. However, it won't provide all the information you need to craft commercial-standard mixes reliably, which is why supplementary monitoring systems are so important.
- The most powerful additional monitoring system for the small studio is exemplified by the classic Auratone 5C Super Sound Cube, a small, unported, single-driver speaker working in mono. This type of speaker highlights the crucial midrange frequencies that reach the largest number of listeners, and it is much less susceptible to the comb filtering and acoustics problems that compromise many affordable nearfield systems. Listening to your mix from a single speaker not only confirms the vital mono-compatibility of your production, but it also makes mix balance judgments far more reliable, especially for important sounds at the center of your stereo image.
- Headphone monitoring is vital to check how your stereo picture and mix balance reach a large sector of the listening public. By isolating you from your listening environment, headphones can also identify low-level technical anomalies and processing side effects, as well as indicating how some room-acoustics problems may be affecting your nearfield listening. Although

comparatively inexpensive headphones will cover all these bases adequately, a top-of-the-range set is highly recommended, because it will usually increase both productivity and mix quality for most small-studio operators, especially for those unable to set up a respectable nearfield system as discussed in Chapter 1.

- Cheap, low-quality "grotbox" speakers also deserve a place in the mixing studio, as they give you an idea of what to expect in worst-case listening scenarios. Use a pair of such speakers placed very close to each other (or a small all-in-one stereo boombox), and then listen off-axis from a reasonable distance to get the "best" results. Although grotbox monitoring is important, don't obsess about it. You should only need one small system to do this job if you're using your other monitors sensibly.

Assignment

- Get hold of a proper Auratone-substitute of some kind, and set up a convenient method of listening to it in mono so that you get into the habit of using it that way.
- Buy a pair of good studio headphones, if possible something at the top of the range so that you can do meaningful mix work when speakers aren't an option.
- Find some suitable grotbox speakers.

www.cambridge-mt.com/ms-ch2.htm

CHAPTER 3

Low-End Damage Limitation

It's the low end of the audio spectrum that presents the toughest mixing challenges when you're working under budgetary constraints, and in my experience it's this aspect of small-studio mixes that most often falls short when compared with professional productions. As I see it, the reason for this trend is that the only real low-frequency monitoring in a typical small studio comes from its nearfields, yet the majority of the small nearfield systems I've encountered in use suffer unacceptably from low-end resonance problems as a result of ported speaker designs and/or room modes. Although there are plenty of ways to reduce resonance effects to workable levels (as discussed in Chapter 1), the sad fact is that few owners of small studios actually invest the necessary money and effort into this area of their listening environment to achieve passably accurate results.

Even if the low end of your monitoring leaves much to be desired or you're forced to work mostly on headphones, all is by no means lost—as long as you're willing to learn a few special mixing and monitoring techniques.

The glib response would be that these guys only have themselves to blame, but many people's budgets are too small to accommodate even relatively affordable monitor-setup and acoustic-treatment measures. There is also a legion of small-studio users who aren't actually at liberty to dictate their gear choice or acoustic setup for mixing purposes: students using college facilities, budding media composers camping out in a corner of a shared living room, or engineers working on location, to give just a few common examples. So it's fortunate that all is by no means lost, even if the low end of your monitoring leaves much to be desired or you're forced to work mostly on headphones. As long as you're willing to learn a few special mixing and monitoring techniques, there are ways to work around the worst low-frequency pitfalls effectively nonetheless and thereby achieve commercial-grade results with some reliability. Indeed, given that low-end problems pop up on almost all small systems to some degree, there's a lot that mixing and monitoring tricks can offer, even to those blessed with respectable nearfield playback.

Mixing Secrets for the Small Studio.

3.1 COPING WITH CHEAP PORTED SPEAKERS

My first set of tips is to help those engineers who find themselves lumbered with having to mix through cheap ported monitors for whatever reason. First, it pays to be aware of where the port's resonant frequency is located, because this knowledge can help you to correctly identify obvious resonances in that region as speaker-design artifacts rather than mix problems. You can also make a note of the pitch of the resonant frequency, which will give you an idea of which bass notes are most likely to suffer irregularity on account of the porting.

You may be able to find out a given speaker's porting frequency from the manufacturer's product specification sheets, but failing that it's straightforward enough to investigate for yourself using the LFSineTones audio file (*www.cambridge-mt.com/ms-ch3.htm*). If you play back this file through your monitor system at a medium volume, you should clearly be able to see the "motion blur" of the woofer's cone as it vibrates back and forth, which makes it easy to tell how wide the woofer's excursions are—it usually helps to look at the speaker cone from the side. As the tones begin their march up the frequency response, a ported monitor's woofer will start off showing fairly wide excursions for the lowest frequencies, but these movements will slowly narrow as the tones approach the speaker's porting frequency. The tone closest to the porting frequency will give the narrowest cone excursions, following which the movements will begin to widen out again as the tones continue on their way up the spectrum. Once you know which tone gives the smallest cone excursion, you can easily refer to Table 1.1 (page 23) to find out the porting frequency, both as a figure in Hertz and as a note pitch.

FIGURE 3.1
Blocking your speaker ports with dusters or bits of foam can give you a useful extra perspective on the low end of your mix.

The second tactic that can help you deal with ported monitors is to block their ports, thereby defeating some of the resonance side effects. Almost any material or foam can be pressed into service here as long as it impedes the flow of air in and out of the port opening—a rolled up sock will normally do a grand job. (What do you mean, "wash it first"? You'll destroy the vintage sound!) Although this will certainly make the speaker significantly lighter on low end, the bass reproduction that you do get should actually be more usable for mix-balancing purposes. Bear in mind, though, that there may well be other disadvantages of blocking the port in terms of frequency response ripples and increased distortion (the speaker was designed to be ported, after all), so though you should be able to judge bass balances more easily with socks stuffed in the port holes, other mixing decisions may work better without them.

3.2 AVERAGING THE ROOM

In Chapter 2 we noted that a decent pair of headphones can bail you out to some extent when room acoustics are undermining the fidelity of your main nearfield monitor system. However, their low-frequency abilities in this department are limited by the size of their drivers, so you still have to rely on your speakers there, which means finding ways to work around any room resonances that are confusing the low-end balance.

Although one of the problems with room resonances is that they vary as you move your monitoring position, this is also the key to one useful workaround: if you make a point of listening to the bass response from several different locations, then it's actually possible to average your impressions mentally to some extent. In a room with resonance problems, you'll find that all aspects of the bass balance will vary as you move around. If a particular bass note, for example, appears to be too loud in some places but too quiet in others, then you can hazard a guess that the overall level may be in the right ballpark, whereas a note that remains too loud all around the room probably needs reining in.

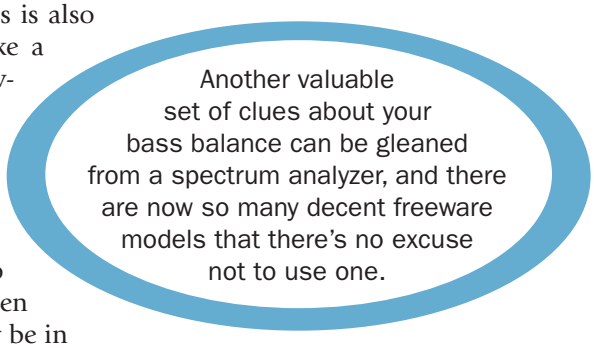

If you want to get the best out of "averaging the room" like this, then jot down your reactions on paper as you move around. This technique can really help clarify your thoughts, especially where there are multiple bass instruments contributing to the mix—kick drum and bass guitar, for example. You'll also find that your proficiency with this trick will improve as you build up some experience of how specific areas of your monitoring room sound. Indeed, it's not unusual to discover a couple of locations in any given room where the bass response as a whole is a lot more reliable than in your stereo sweet spot. A valuable additional perspective can also be gained by listening to your mix through the doorway from an adjoining room, as the different acoustic space will filter the sound through a different set of room modes—Allen Sides,[1] Joe Chiccarelli,[2] and George Massenburg[3] are just three of the numerous well-known engineers who mention this little dodge.

No matter how methodical you are about it, though, averaging the room will never be an exact science. It can provide a lot of useful extra clues about the low-end balance, but you should nonetheless beware of basing drastic mixing decisions on conclusions drawn solely from this source.

3.3 SPECTRUM ANALYSIS AND METERING

Another set of clues about your bass balance (and indeed your mix's overall tonality) can be gleaned from a spectrum analyzer, and there are now so many decent freeware models that there's no excuse not to use one—RN

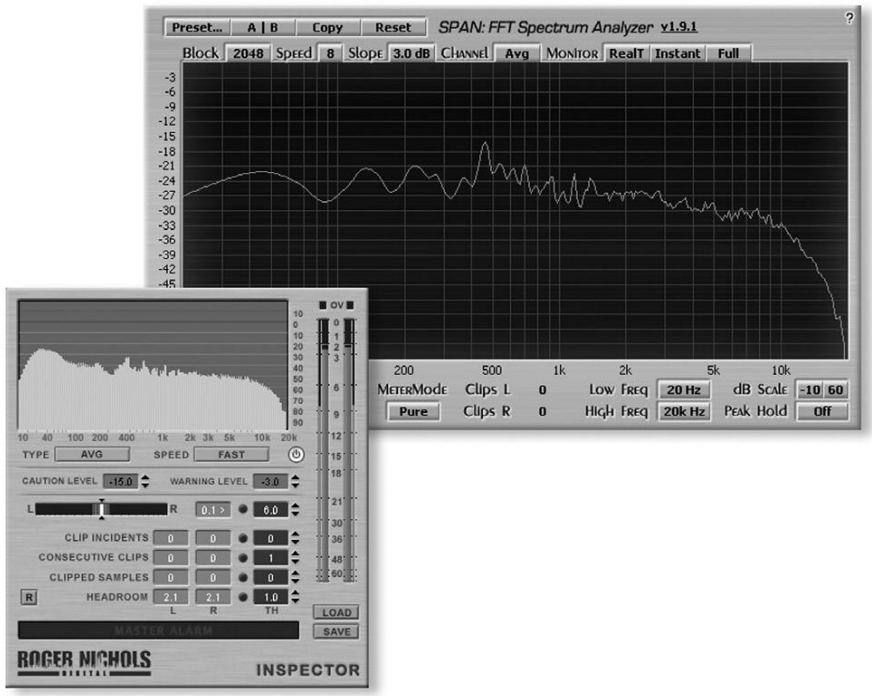

FIGURE 3.2
Spectrum analyzers have a lot to offer the small-studio mix engineer. Two good freeware models are RN Digital Inspector and Voxengo SPAN.

Digital's Inspector and Voxengo's SPAN are just two great examples. Engineers Joe Chiccarelli[4] and Eric Rosse[5] have both mentioned using spectrum analysis. "I put [a spectrum analyzer] across my stereo buss that lets me know when the bottom end is right," says Chiccarelli. "I'm mainly looking at the balance of the octaves on the bottom end, like if there's too much 30Hz but not enough 50Hz or 80Hz. When you go to a lot of rooms, that's where the problem areas of the control room are." You should try to find an analyzer that provides good resolution in the frequency domain (Inspector has about the minimum resolution I'd recommend in this regard), and it helps to have some control over the metering time response, so that you can switch between slower averaged meter ballistics (which will be better for overall level judgments) and faster peak metering (which will track things like drum hits more closely).

The most important thing to understand about spectrum analyzers, however, is that each manufacturer's implementation will present the frequency information in a slightly different way. This means that you can only really begin interpreting a spectrum analyzer's display usefully once you've built some experience of how your particular model responds to commercial material. There's no sense in trying to line up all the bars on the graph with your school

ruler if that's not what the target market expects. Remember also that spectrum analyzers can evaluate individual tracks as well as whole mixes, and they're particularly handy for highlighting if the internal frequency components of a bass part are shifting from note to note—a problem that all too often scuppers small-studio balances.

Don't ignore your recording system's normal level meters either, as these can reveal some undesirable level irregularities on individual instruments. Again, you'll get more information if you use a meter that can show both average and peak levels, and fortunately such meters are commonly built into most digital audio workstation (DAW) software these days. (If you don't have one on your own platform, then there are some good dedicated freeware models such as Sonalksis FreeG or Tischmeyer Technology's TT Dynamic Range Meter. RN Digital's Inspector also includes one.)

3.4 WATCH THOSE CONES!

A final visual indicator of potential low-frequency concerns is your speaker cones themselves. "You can sometimes tell how much low end you have on an NS10 from the way the woofer moves," says Manny Marroquin.[6] As I mentioned when talking about finding a speaker's porting frequency, the woofer cone excursions become visible at lower frequencies and remain so even at frequencies that are too low either to be delivered effectively from a typical speaker cone or indeed to be perceived by human hearing. Particularly on low-budget productions, the buildup of inaudible subsonic energy is a real hazard, because it can interfere with the correct operation of your channel mix processors, prevent your finished mix from achieving a commercially competitive loudness level, create unpleasant distortion side effects on consumer playback devices, and make the mix translate less well when broadcast.

FIGURE 3.3
Excellent high-resolution peak/average metering is freely available in the form of the Sonalksis FreeG and Tischmeyer Technology TT Dynamic Range plug-ins.

"Many people want the bass to be really loud," says Dave Pensado. "But if it's too loud the apparent level of your mix will be lower on the radio. If you put in too much bass, every time the [kick] hits the vocal level sounds like its dropping 3dB."[7] Just keeping an eye on your woofer excursions can be a valuable safety check in this respect, because a lot of sounds that don't seem to have low end to them can still incorporate heaps of subsonics. "You've got to make sure that you're not adding sub-sonic stuff," says Chuck Ainlay. "If I see a lot of excursion on the woofer, then I'll start filtering something. A lot of times it exists in a bass guitar, which can go quite low, but what you're seeing there is a sub-sonic harmonic. That can be filtered out without hindering the sound of the bass at all."[8]

Other common culprits are vocal "p," "b," and "w" sounds, which can generate subsonic thumps by dint of their hitting the microphone's diaphragm with a blast of air. A performer's movements

can also create drafts around the microphone, which have a similar effect—close-miked acoustic guitars can be prone to this, for instance. The vibrations of traffic rumble or foot tapping can easily arrive at a microphone via its stand too. And the issues aren't just restricted to mic signals. A lot of snare drum samples have unwanted elements of low-frequency rumble to them because they've been chopped out of a full mix where the tail of a bass note or kick drum hit is present. Aging circuit components in some electronic instruments can sometimes allow an element of direct current (DC) to reach their outputs. This is the ultimate subsonic signal, as it's effectively at 0Hz and offsets the entire waveform, increasing your mix's peak signal levels without any increase in perceived volume.

Even with the low-frequency energy that you actually want in your mix, cone movements can occasionally be a useful indicator. For example, it's not uncommon with both live and sampled kick drums for the lowest-frequency components to be delayed compared to the higher-frequency attack noise, and this often isn't a good thing when you're working on hard-hitting urban, rock, or electronic music—it just sounds less punchy. You may not be able to hear the lowest frequencies of a kick drum on a set of small nearfields, but if the onset of the corresponding cone excursions seems visually to be lagging behind the beat, then it can be a tip-off that your choice of sound may not deliver the goods on bigger systems. Many bass-synth presets also feature an overblown fundamental frequency, and if you're not careful with your MIDI programming, this fundamental can wind up wandering uselessly underneath the audible spectrum. A franticly flapping woofer cone (or hurricane-force gusts from a speaker port) can flag up such problems even though you can't hear them. For an illustration, check out my two ConeFlapper audio examples (but at reasonably low volume level to avoid any damage to your speakers): ConeFlapperOut gives a section of an R&B-style backing track with tightly controlled low end; ConeFlapperIn, by contrast, has a strong subsonic element to the kick drum, which eats up an extra 3dB of mix headroom and will flap your woofer around like crazy, but otherwise creates no significant change in the sound at all (*www.cambridge-mt.com/ms-ch3.htm*).

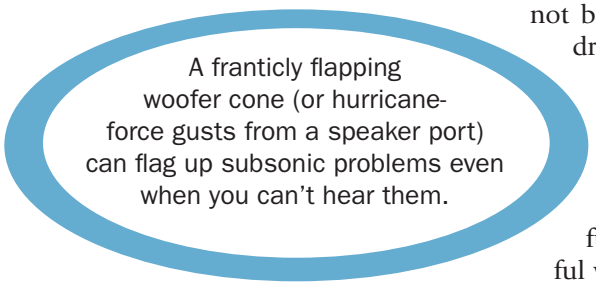

Once more, though, I want to sound a note of caution, because you should avoid trying to conclude too much about your mix from woofer wobbles. Some engineers, for example, suggest that resting a fingertip on a small speaker's woofer cone allows you to feel imbalance and unevenness in the bottom octaves of the mix. In my view, the evidence of cone excursions isn't actually much help with these kinds of mix tasks, especially if you're using ported monitors, which reduce their excursions around the porting frequency as mentioned earlier. Even with unported monitors, cone excursions still tend to increase for lower frequencies, because that allows some compensation for the way that

the woofer's limited size becomes less efficient at transferring lower-frequency vibrations to the air. As with all monitors and monitoring techniques, you have to take from your woofer movements only the information that's of practical use and be merciless in disregarding anything that might mislead you.

3.5 PREEMPTIVE STRIKES AT THE LOW END

Despite what all these workarounds have to offer users of small studios, there will inevitably be some unwelcome degree of guesswork involved when crafting the low end of a mix unless you have at least some access to a reasonably well-behaved, full-range nearfield system. Faced with this uncertainty, then, the canny engineer will employ a certain amount of preemptive processing to avoid any low-end problems that the available monitoring can't adequately detect, and will also deliberately craft the final mix so that it responds well to mastering-style adjustments should aspects of the low-end balance prove, with hindsight, to have been misjudged.

Simplify the Problem, Simplify the Solution

Chief among these preemptive strategies is to restrict the field of battle, as far as low frequencies are concerned. In the first instance, this simply means high-pass filtering every track in your mix to remove any unwanted low frequencies. (I'll deal with the specifics of this in Chapter 8.) "[I use] a simple high-pass filter... on almost everything," says Phil Tan, "because, apart from the kick drum and the bass, there's generally not much going on below 120 to 150Hz. I have always found filtering below this cleans up unnecessary muddy low-level things."[9] Just because an instrument is supposed to be contributing low frequencies to your mix, that doesn't mean you shouldn't high-pass filter it, either, because even if the filter barely grazes the audible frequency response, it will still stop troublesome subsonic rubbish from eating away at your final mixdown's headroom. As Serge Tsai observes, "In general I like to take off rumble from the room things were recorded in. I like to keep my low end clean."[10]

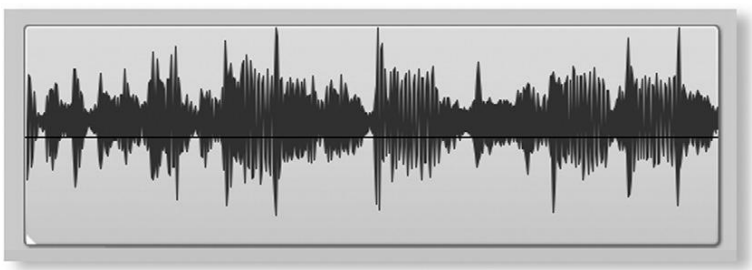

FIGURE 3.4
The effect of DC (0Hz) on a mix file's waveform. Notice how the positive waveform peaks are clipping, even though the negative waveform peaks still have headroom to spare.

Beyond this general trash removal, there are advantages for the small-studio engineer in deliberately reducing the number of tracks that include low-end information. "You've got to remember," advises Jack Douglas, "that the stuff that's going to take up the most room in your mix is on the bottom end. If you just let the bass take up that space, you can get out a lot of the low stuff on other tracks—up to around 160Hz—and it will still sound massive."[11] So if, for example, your main bass part is an electric bass, but there's also significant sub-100Hz information coming from your electric-guitar, synth-pad, piano, and Hammond-organ tracks as well, then it's not a bad idea to minimize the low-end contributions of all these secondary parts to reduce the overlap. This not only means that you can fade up the main bass part more within the available headroom, but the sub-100Hz region will also become much easier to control, because you can concentrate your low-frequency processing on just the bass guitar track. You might even split off the sub-100Hz frequency components for separate processing, metering their levels to ensure that they remain rock solid in the mix balance. Or perhaps you might decide to replace those frequencies completely with a dedicated subbass synthesizer part (a common hip-hop and R&B trick), using your MIDI and synth-programming skills to dictate the level of low end with absolute precision.

Restricting any serious low end in your mix to the smallest number of tracks possible has a couple of other important advantages too. First, it helps the monitoring and metering workarounds described earlier to be more effective in practice, simply because what's going on in the bottom octaves is subsequently less complicated to unravel. Second, should you discover post-mixdown that the weightiness of your main bass instruments is out of line with the market competition, then you can usually do a lot more to correct this frequency problem using mastering-style processing, without destroying the tone of other instruments or compromising the clarity of the mix as a whole.

Clearly, hedging your bets at the low end like this must inevitably cramp your sonic style to some extent, and you may sacrifice a little low-end warmth and nuance by working in this way. However, you can't expect to have your cake and eat it too. If you don't have commercial-grade low-frequency monitoring, but nonetheless want the kind of clean, consistent, and powerful low end you hear on commercial records, then you've got to consider the demise of a few comparatively minor sonic niceties a pretty meager price to pay.

CUT TO THE CHASE

- The single biggest mixing challenge in the small studio is getting the low end right, partly because the monitoring tool best suited to tackling it is also the most expensive and the most complicated to set up effectively: a full-range nearfield system. However, for those who have only limited access to a decent nearfield monitoring environment or those without the means to remedy the shortcomings of a compromised loudspeaker system,

there are a number of workarounds that allow respectable low-end balances to be achieved nonetheless.

- If you have to use ported monitors, then knowing the port's resonant frequency can help you mentally compensate for some of its side effects. Blocking the speaker ports can also improve your ability to make reliable decisions about the low end, although this may detract from the speaker's performance in other important areas.
- If you can't adequately tackle resonant modes in your listening room (and few small studios make sufficient effort to do so), then you'll make better low-end mixing decisions if you "average the room," comparing the mix's bass response from different locations inside and outside your studio monitoring environment.
- Good spectrum-analysis and metering software can be very helpful to your understanding of the low frequencies in your mix. However, you need to spend time acclimatizing yourself to the way any particular meter responds before you can get the best from it.
- The visible movements of your speaker cones can warn you of undesirable subsonic information in your mix and can reveal some areas for concern with different bass instruments. Beware of reading too much into cone movements, though, because they can give misleading impressions of low-end balance.
- If you're in any doubt about what's going on at the low end of your mix, then try to simplify the problem by eliminating unwanted low-frequency information and reducing the number of tracks that carry significant sub-100Hz information. This will allow you to maintain better control over the low-frequency spectrum; the monitoring workarounds already suggested will assist you more effectively; and post-mixdown mastering processing will be better able to remedy any low-end balance misjudgments you may inadvertently have made.

Assignment

- If you're using ported monitors, work out their porting frequency and make a note of the pitch it corresponds to.
- Find yourself a level meter that shows both peak and average levels, and also a high-resolution spectrum analyzer. Use them while mixing and referencing so that you get to know how they respond in practice.

www.cambridge-mt.com/ms-ch3.htm

This page intentionally left blank

CHAPTER 4

From Subjective Impressions to Objective Results

Monitoring is usually the biggest stumbling block in the path of most small-studio operators, but even the most laser-guided listening system can't mix your tracks for you. It'll tell you where you are, but it's still your job to navigate to where you want to be. The basic difficulty is that listening is an inherently subjective activity. The same "mix in progress" will usually elicit a different reaction from every different listener (even from the same listener after a heavy lunch!), and different engineers will all suggest different processing tweaks. Unless small-studio engineers can make their mixing decisions more objective, they will always lack the confidence to direct mixing resources efficiently and successfully.

Now it would be madness to assert that entirely subjective mix decisions have no place in the mixing process. Of course they do, and so they should. That's how people will know that a given mix is yours rather than mine, and it's this kind of contrast that's essential for the development of new and exciting music. But no book can teach you creativity, so I'm afraid I can't help you with the truly subjective decisions. Not that this actually matters a great deal, because most small-studio users have more ideas in their heads than they know what to do with, so even if I could give you a magic bottle of inspiration, I'm not sure you'd actually be any better off for drinking it! More to the point, it's the inability to make more objective decisions that really holds back most small-studio productions. You can create the vocal sound that launched a thousand ships, but it won't launch a pedalo unless you can make it fit properly in your mix.

This chapter takes an in-depth look at techniques you can use to make the kinds of down-to-earth decisions that really matter to commercial mixing. Can I hear the lyrics well enough? Do I have too much sub-bass? Where should I pan my overhead mics? How much delay should I use, and how much reverb? Answering these and similar questions is the bulk of what mixing work is about, and unless you can answer them appropriately, even a truly inspired mix will sound flat when set against commercial tracks from your target market.

Mixing Secrets for the Small Studio.

4.1 FIGHTING YOUR OWN EARS

Ears are pretty important if you want to hear anything, but they're a nuisance when you're trying to stay objective about your mix. This is because the human auditory system doesn't just transmit sonic vibrations straight from the air to your consciousness; it not only colors the raw sensory input through its own nonlinear response, but it also constantly adapts itself in response to incoming sounds—partly in order to extract the maximum information from them, and partly just to shield its sensitive physical components from damage. Although the fact that our hearing works like this is helpful in everyday life, it actually works against you when mixing, casting a long shadow of doubt over every balancing and processing decision you make. If your hearing system is doing its best behind the scenes to balance and clarify your perception of the mix, rather than consistently presenting you with the warts-and-all reality, then it's all too easy to be suckered into signing off a second-rate mix as a finished product. The key to eluding this trap is to understand a bit about how our physiology skews what we hear, so that you can work around or compensate for the perceptual quirks when evaluating your mix.

Shock Tactics

One of the main problems is that the ear is very good at compensating for tonal imbalances. In one respect this is useful, because it means that tonal differences between brands of monitor speaker don't make as big a difference to their usefulness at mixdown as you might expect. Once you've acclimatized to a particular set of speakers, your ear will factor out their unique tonal characteristics to a certain degree. But for every silver lining there's a cloud: even if your mix has an obviously wonky frequency balance, it only takes a few seconds for your hearing system to start compensating for that, thereby hiding the problem from your consciousness.

If your hearing system is doing its best behind the scenes to balance and clarify your perception of the mix, then it's all too easy to be suckered into signing off a second-rate mix as finished product.

How can we combat the ear's fickleness? Well, one good tactic is to switch between monitoring systems fairly frequently, because this instantly changes the tonal frame of reference and hence offers a few precious moments of clarity before the hearing system has a chance to recover from the shock and recalibrate itself. "After listening to all sorts of frequencies for a long period," says Russ Elevado, "the Auratone will just put everything flat and my ears kind of just recalibrate themselves."[1] Given that each of your monitoring systems offers a different set of strengths and weakness as far as analyzing your mix is concerned, switching between them is already a good idea anyway, so this is one more reason to make a habit of it. It's no accident that the monitor-selection buttons on most large-format professional studio consoles are usually among the most worn looking. "I can easily switch between 15

different speakers," says Cenzo Townshend, for instance. "I couldn't just mix on NS10s or KRK 9000s… you get used to them too much."[2] In fact, because switching monitors is so useful while mixing, I usually recommend that small-studio owners without a suitable monitoring section in their mixer get hold of a dedicated hardware monitor controller for the purpose, because it saves a great deal of time in the long run if you can hop between your monitor systems at the flick of a switch—and it makes getting into the routine of doing it easier too. There are some extremely affordable models on the market, but you should bear in mind that everything you hear will pass through your monitor controller, so it's perhaps a false economy to pinch pennies here.

FIGURE 4.1
An excellent tool for recalibrating your ears. It also makes hot drinks.

A second important shock tactic is taking breaks. According to Bob Clearmountain: "Don't just keep plugging away at something. When you start to question everything that you're doing and you suddenly don't feel sure what's going on, stop for a while… I go out with my dog or take a 15-minute catnap, then come back and have a cup of coffee and it's much more obvious if something is wrong."[3] "The tougher the mix, the more breaks you should take," says Russ Elevado. "If it's a dense song, and there's lots of frequencies going on, I'll do 20 minutes, then take a 20 minute or more break. On a difficult song, you can be more productive with more breaks… It's definitely good to pace yourself and know that you're not wasting time when you're taking a rest. You can get too focused on the details when you're working for a long stretch. Once you take a nice break, you can come back and are more able to look at the overall picture."[4]

Taking a break doesn't just mean sagging back in your studio chair and scoffing pizza, either—leave your studio room and do something to remind your ears what the real world sounds like. If anything you hear in your studio setup is a better reality check for your ears than the sound of your own kettle boiling, then you seriously need to get out more! The longer you work on a track, the longer your breaks may need to be. As Bill Bottrell explains: "Objectivity is everybody's big problem, especially if you've worked a long time on a song… But, if something ends up sounding wrong, you will hear that it's wrong if you get away from it for five days and then hear it again."[5]

Breaks and Hearing Fatigue

The benefits of taking breaks, though, extend beyond basic shock tactics. Whatever volume level you mix at and however smooth-sounding your music,

a few hours of uninterrupted mixing will still tire out your ears, and your sensitivity to high frequencies in particular will suffer. So if you want to be able to mix all day the way the pros can, then breaks are essential to keep your ears from flagging. Even then it pays to be careful about what you commit to toward the end of a long session. After one grueling (and not particularly successful) mixing project, I came up with a basic rule of thumb that has served me well ever since: Never finalize a mix after supper. I find that by that time of the day my ears simply aren't objective enough to do the job properly. If at all possible, try to "sleep on the mix" before finishing it. "When you come back fresh," says Joe Chiccarelli, "there are always a couple of obvious little things that you've overlooked."[6]

Engineers such as Bob Clearmountain,[7] Mike Clink,[8] Russ Elevado,[9] and Thom Panunzio[10] have all echoed this sentiment in interviews, whereas Joe Zook offers this cautionary tale about a hit single he mixed for One Republic: "When I first mixed 'Stop & Stare,' I made the drums and the bass massive. It became this huge grunge thing, and it just didn't sound unique or fresh. I'd started at five or six in the evening and I went home late at night being really proud of myself. But when I came back in the next morning, it was like, 'F**k, where did the song go?' There were all these great sounds, but it didn't feel right… I'd become too focused on all the unimportant aspects of the engineering, rather than the music itself. So I started again."[11]

Monitoring Level

The simple issue of how loud you listen while mixing also needs some thought if you're going to get the best results. For one thing (and to get the preaching out of the way as quickly as possible), it's just plain daft to listen at thundering volumes for long periods because there's a real risk of hearing damage. If you reach the end of a day's mixing and find that your hearing's noticeably dulled or there's a high-pitched ringing in your ears, then you've probably damaged your hearing to some degree. Every time you do that, you're eating away at your hearing bit by bit, and there's no way to bring the sensitivity back once it's gone. Most of the highest-profile mix engineers actually spend the majority of their time mixing at low volume. Take Allen Sides, for instance. He stated in an interview that he only listens loud for maybe 20 to 30 seconds at a time and does his most critical balancing at a low enough level that he can easily conduct a normal conversation during playback.[12] So think like a professional; take care of your ears!

It's just plain daft to listen at thundering volumes, because there's a real risk of hearing damage. Most of the highest-profile mix engineers actually spend the majority of their time mixing at quite low volume.

"Be safety conscious when you go to shows, and monitor at reasonable volumes," advises Kevin Killen. "Do not be afraid to protect your most valuable commodity."[13] Take special care when listening on headphones, because the unnatural listening experience they create makes it easy to monitor dangerously loud. And be aware that alcohol and drugs can tamper with the physiological protection systems that are built into your hearing system, so be doubly cautious if you're under the influence. (Yeah, yeah, of course you don't.)

There's more to setting your monitoring level than avoiding concussion, though, primarily because these pesky ears of ours present us with a different perceived frequency balance depending on how loud we listen. If you fancy researching the psychoacoustics of this effect, then graphs such as the well-known equal-loudness contours shown in Figure 4.2 will clue you up on the details, but for mixing purposes you only really need to understand this quick-and-dirty generalization: you'll hear more of the frequency extremes as you crank up the volume. That's not the only factor at work, though, because our ears actually reduce the dynamic range of what we're hearing at high volumes as a safety measure too, and background noise effectively recedes as monitoring levels increase.

The most immediate ramification of all this is that you need to understand how your production responds to being played at different volumes. "I switch my monitor levels all the time," remarks Bob Clearmountain. "I tell people I work with 'Just give that monitor pot a spin any time you like.'"[14] As Chuck Ainlay clarifies: "I don't see any way of identifying if the whole mix is really holding up unless you crank it. I don't know how you can check the bottom end and whether the vocals are poking out too much without turning it up quite loud. There are also things that you'll miss by listening loud, so you have to work at low levels too."[15]

However, although listening at a wide variety of volume levels will help you keep perspective on how your mix translates across all systems in general, it's equally important for the mix engineer to take an active decision as to which volume levels are of the greatest importance for the musical style in question. For example, a teen pop/rock mix might require serious high-frequency enhancement to retain its clarity in low-level home/car/office listening situations, whereas the same degree of fizz coming full-throttle over a club sound system would all but tear your ears off! Furthermore, it's important to realize that the inherent excitement of loud listening can make even a limp mix seem powerful. "I do most of my mixing at pretty low level," says Clearmountain. "I get a better perspective that way and I can work longer. If I can get a sense of power out of a mix when it's quiet, then I know that's still gonna be there when I turn it up."[16] Chris Lord-Alge agrees: "I listen really quietly. When you've been doing this job for a while, you tend to listen at such a low volume that in order for it to hit you in the face, you have to really push it."[17]

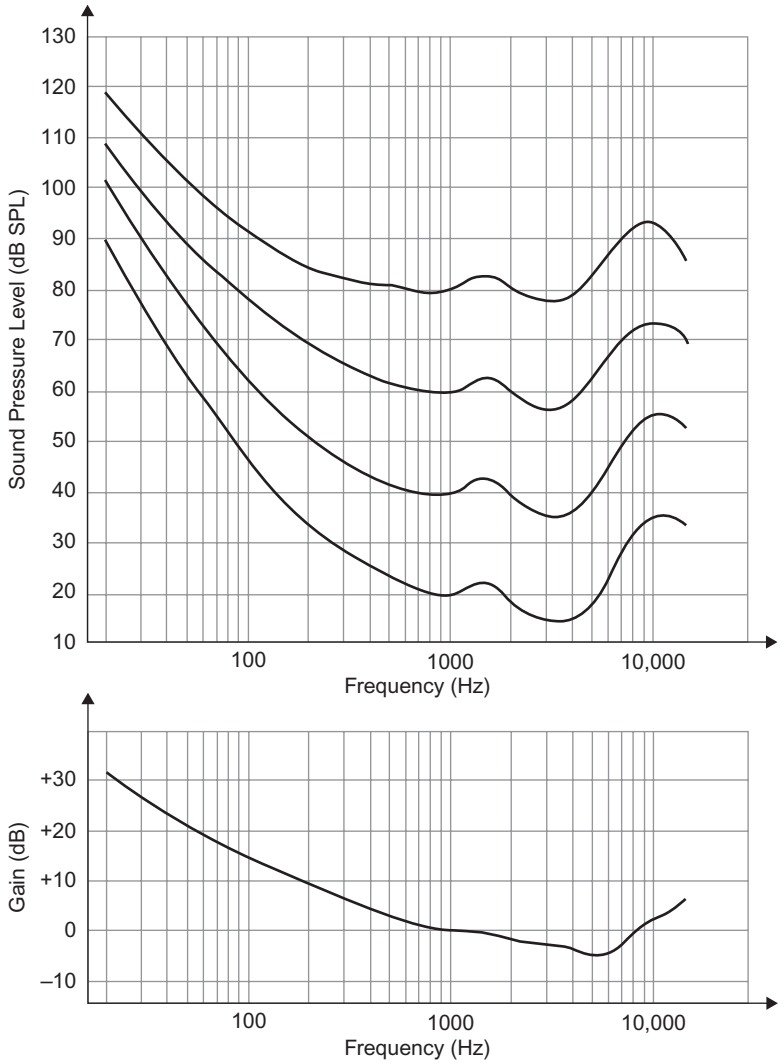

FIGURE 4.2
The top graph here shows four equal-loudness contours. One of the things they show is that our sensitivity to frequency extremes increases with listening level—for instance, moving from the bottom to the top contour (in other words turning the volume up by 60dB) has roughly the same effect on the perceived tonality as the EQ plot shown at the bottom.

A Jury of Your Peers

Nervous about playing your mixes to other people? Well, grow up and get over it! I'm not trying to be macho for the sake of it here, it's just that mixing on a professional level is about taking responsibility for your output, and that means playing it unflinchingly for all and sundry. By the time you've worked your way through the mixing process, your own ears will probably be the least

objective on the planet about the outcome, which means that it's easy to miss the wood for the trees. There are few reality checks to match another person's dispassionate first impression, so any aspiring mix engineer with even half an ounce of sense should seek out constructive criticism wherever possible.

"You've got to be very careful about losing your objectivity," warns Alan Moulder, recalling his mixing work for Smashing Pumpkins: "It was very handy to have [producer] Flood and [lead singer Billy Corgan] there. They could come in and say, 'That sounds great,' and I'd say, 'Well, so-and-so's not as good,' and one of them would say, 'No, it's fine. I prefer it.' That meant I could just lose myself in the technical world for a while, because these other guys hadn't lost their objectivity."[18]

Alan Parsons also advises against going it alone: "I don't think any single engineer, artist, or producer ever achieved anything by working on his own. You learn from other people, and you achieve great things by interacting with other people."[19] Play the mix to your music technology lecturer and watch for winces. Bribe the local DJ to sneak it into his set and find out if it clears the floor. Play it to your pot plant and see if it wilts. The more you get your work out into the public domain, the more informed you'll be about the merits of different mixing techniques, and the quicker your skills will progress.

Clearly you'll want to take some people's comments with a large pinch of salt, because every individual has his or her own opinion about what constitutes good music. Few death metal fans are likely to heap praise on a plastic Euro-trance single, for instance, however stellar the production values. You have to learn how to interpret each criticism, filtering out what you decide is irrelevant and using the remainder to improve your sound. If you can get good at this when you're starting out, it'll stand you in excellent stead for dealing with the comments and revision requests of paying clients further down the line.

Here are a couple of tips that can help with interpreting other people's opinions of your mix. First, the most useful advice tends to come from those listeners who really connect with the sound you're aiming for and seem to like the general approach you've taken with your mix. When that kind of listener says that some aspect of the mix doesn't feel right, then it's usually wise to pay heed. Second, it's not uncommon for people's comments to be apparently nonsensical or contradictory, especially if you're dealing with nonmusicians. Although this can be frustrating, it's fair enough. Music isn't an easy thing to describe at the best of times, and much of the language people use to do so is vague. To get around this problem, communicate using music itself. In other words, ask the critic in question to identify commercial mixes that exemplify the kind of attribute that he/she feels your mix is lacking. Let's say, for example, that one of your listeners wants a "more powerful snare." Get the person to name-check a few professional releases with such a snare sound, and it suddenly becomes easier to work out whether your fader, EQ, compression, or reverb settings need adjusting—or whether that person is, in fact, spouting utter tripe.

4.2 THE ART OF MIX REFERENCING

Although each of the monitoring tactics I've discussed so far in this chapter is important in its own right, I've reserved special attention for the most powerful technique of all: mix referencing. George Massenburg neatly encapsulates what this process is all about: "I wish guys would pick up several of the top 100 CDs—the ones they're trying to emulate—and listen to them, compared with their work. And then tell themselves that their ears are not deceiving them!"[20] Particularly for those starting out, mix referencing against commercial CDs is a dispiriting process, so it's little wonder that so many small-studio users shy away from it. However, this painful trial by fire is also quite simply the best mixing tutor in the world, and the fastest track to commercial-sounding productions.

The fundamental point of mix referencing is that it delivers what your own hearing system and experience cannot: objective decisions. Because we humans are only really equipped to make relative judgments about sound, the only way you can anchor those judgments in reality is by comparing your mix against existing commercial-quality productions. Outside the realms of high-spec professional rooms, mix referencing takes on another role too, because it allows you to compensate somewhat for the skewed listening perspective of unavoidably compromised monitoring systems. In other words, you can start to answer questions like "Do I have too much 100Hz, or are my speakers telling a big fat lie?"

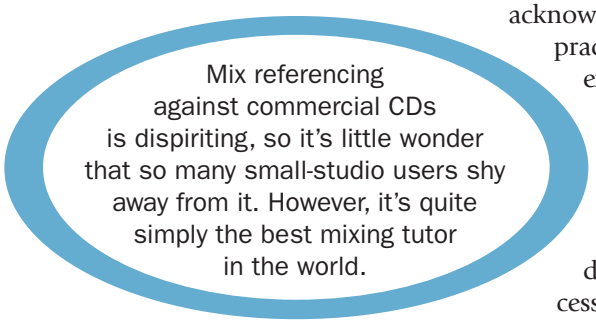

Although every sensible small-studio operator eventually acknowledges the basic concept of mix referencing, in practice most people working on a small scale fail to extract the maximum benefit from the referencing process. This is a real shame, because conscientious referencing is the ultimate bang-per-buck studio tool. It's the best way of closing the gap between amateur and professional sonics, yet it costs very little money. So let's delve more deeply into the mechanics of the referencing process to demonstrate some of the refinements.

Choosing Your Material

Easily the biggest mistake you can make when mix referencing is to use inappropriate comparison material, and the only decent way to avoid this pitfall is to give the selection process a proper slice of your time. This means vetting suitable tracks from your record collection as methodically as you can, on a variety of listening systems, and then comparing them side by side to knock out the lesser contenders. Sounds like a monumental drag, doesn't it? I agree, but let me tell you a little story.

When I first started mixing, I filled up a CD-R for referencing purposes with a couple of dozen commercial tracks, selecting my favorites pretty much from

FIGURE 4.3
Both these records have classic singles on them, but you may find that the mix sonics are better on other tracks from each album.

memory—much as most small-studio users tend to. Not long afterward, I began to notice that other commercial mixes seemed to sound better than my reference material, so I decided to take a more methodical selection approach and systematically trawled through my entire record collection, making an effort to ratify each track preference on at least three listening systems. (Sad, I know, but I'm a bit of a geek.) After much compiling and recompiling of short lists, I discovered that only one track from my original disc remained among my newly minted top 60. By putting my mixes side by side with lower-grade reference material, I'd effectively been compromising the quality of my own output because I'd not had to work as hard to achieve comparable results.

In a lot of cases, I'd instinctively selected my favorite song off an album, rather than the most useful mix reference from a sonic perspective. In other cases, less high-profile records by a given mix engineer actually sounded better than the mega-famous example I'd initially chosen—Andy Wallace's mix of Nirvana's "Smells Like Teen Spirit" made way for his mix of Rage Against The Machine's "Fistful of Steel," for example. The moral of this tale is that your instincts will play tricks on you if you give them half a chance. The whole point of mix references is that they should allow you to foil such perceptual skullduggery, but they'll only do this properly if you factor out the vagaries of your own hearing system during the selection process. I'm not suggesting that everyone take as tortuous a route as I did (it's probably not healthy), but the more ruthless your shootout procedure, the more powerful a tool your mix references are likely to become.

What Makes a Good Reference?

Of course, you can be as anal about selecting reference tracks as you like, but that won't help you if you don't have an idea of what you're listening for. Clearly, your overall judgment of the sound quality is an important factor when selecting mix reference tracks, and it's right that this judgment should be a subjective one to a certain extent—it's your unique idea of what sounds good that will give your mixes their recognizable character, after all. However, it's nonetheless sensible from a commercial perspective to make the effort to understand and evaluate productions that have achieved an elevated status among the professional community in terms of sheer audio quality.

Any engineer you ask will usually be more than happy to offer recommendations in this regard, and Internet forums such as those at *www.soundonsound.com* and *www.gearslutz.com* are regularly host to impassioned "best-sounding album in the Universe ever!!!!" exchanges that can provide food for thought. Another extremely well-regarded source of suggested reference tracks can be

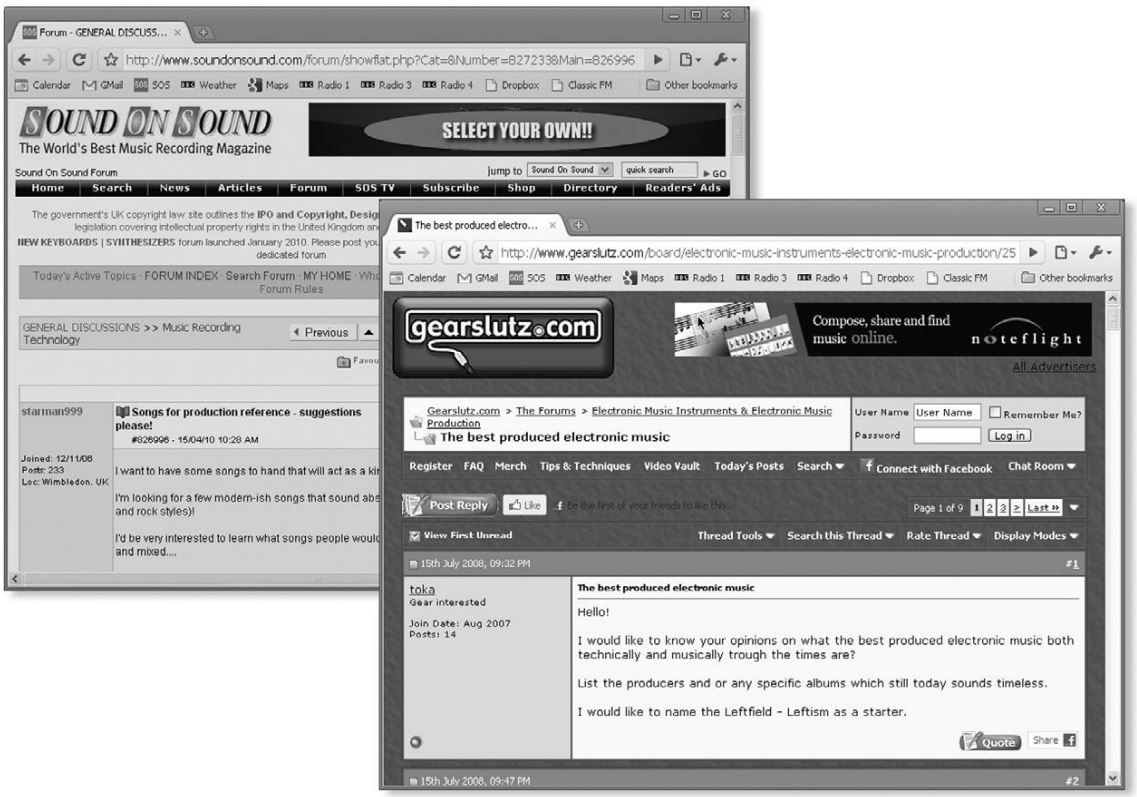

FIGURE 4.4
If you need suggestions for good-sounding records, there are a hundred lists to be found on web forums. Check out *www.soundonsound.com* and *www.gearslutz.com* in particular.

found on the website of top mastering engineer Bob Katz. His "Honor Roll" page at *www.digido.com* lists dozens of albums in all styles. That said, if you browse around his site it quickly becomes clear that Katz is a leading figure in the fight against overloud "hypercompressed" CDs. So without engaging in that whole debate myself here, you do need to bear this in mind. If, for whatever reason, you want the loudness of your mixes to compete with recent commercial records that use heavy-handed mastering processing, you may find the Honor Roll of limited use for referencing purposes.

Beyond just seeking out high-quality sonics, it's good sense to go for reference tracks that relate to the kinds of musical styles you work with, because production values can vary dramatically between styles. "You want to sound contemporary and current," explains Jimmy Douglass, "but you can't know what that is unless you listen to the records that the audience is digging at the moment."[21] There's also the issue of what era of record production you want to emulate. Approaches to the use of compression and reverb, for example, have changed a great deal over time, and tastes continually shift regarding overall mix tonality. However, while you may wish to weed out reference tracks if they don't conform to the genres you work in, it's important to resist the powerful urge to eliminate tracks purely on grounds of musical taste. It's not that quality songs aren't important in a general sense, but it's the sonics of a track that make it a good mix reference, rather than the songwriting. So while I personally remain unconvinced of the musical value of the Pink, Craig David, and Puddle of Mudd tracks on my own reference CD, they've all proven their worth repeatedly for my mixing work.

It's important to resist the powerful urge to eliminate reference tracks purely on grounds of musical taste. It's the sonics of a production that make it a good mix reference, rather than the songwriting. So while I personally remain unconvinced of the musical value of the Pink, Craig David, and Puddle of Mudd tracks on my own reference CD, they've all proven their worth repeatedly for my mixing work.

Another crucial thing to realize is that you shouldn't try to reference your whole mix from a single track, because different tracks will be useful for validating different mix decisions. For example, Skunk Anansie's "Infidelity (Only You)," Sting's "All Four Seasons," and Dr Dre's "Housewife" are all tracks I regularly use to help me mix bass instruments. By contrast, I might use Cascada's "Hold You Hands Up," Gabrielle's "Independence Day," Norah Jones's "Sunrise," or Paolo Nutini's "New Shoes" for exposed vocals. Other tracks can be useful as "endstop markers," such as Natalie Imbruglia's "Torn," where the vocal air and sibilance are pushed too far for my liking. If a vocal in my mix approaches "Torn," then it means that I have to rein in the high end—and that it's probably time to comb the Chihuahua or water the triffids or something.

Reference material is particularly important where the use of room ambience and reverb is concerned, as this is one area of a production that can be tricky to get right. For instance, Outkast's "Bowtie," Solomon Burke's "Other Side of the

Coin," Kanye West's "Breath In, Breathe Out," and Keane's "She Has No Time" all take different approaches in this respect, and having access to all of them among my reference tracks helps me make more suitable use of effects across a range of styles. Mix-buss processing is also easy to overdo without some good examples to compare against, and audience expectations of mix tonality can only really be judged with representative tracks from the chosen target market on hand.

"You can do a mix where all the balances are perfect," says Andy Johns, "but it means shit unless the overall sound—the bottom and the top—is going 'hallo'! If it ain't there, see you later… You have to have something to compare it to, so I'll constantly be playing CDs."[22] Once more, don't neglect the usefulness of "end-stop" examples here. The tonality of Madonna's "Sorry" works like this for me at the high end, for instance, whereas Pussy Cat Dolls' "Taking over the World" and Eminem's "Square Dance" provide a similar service at low frequencies.

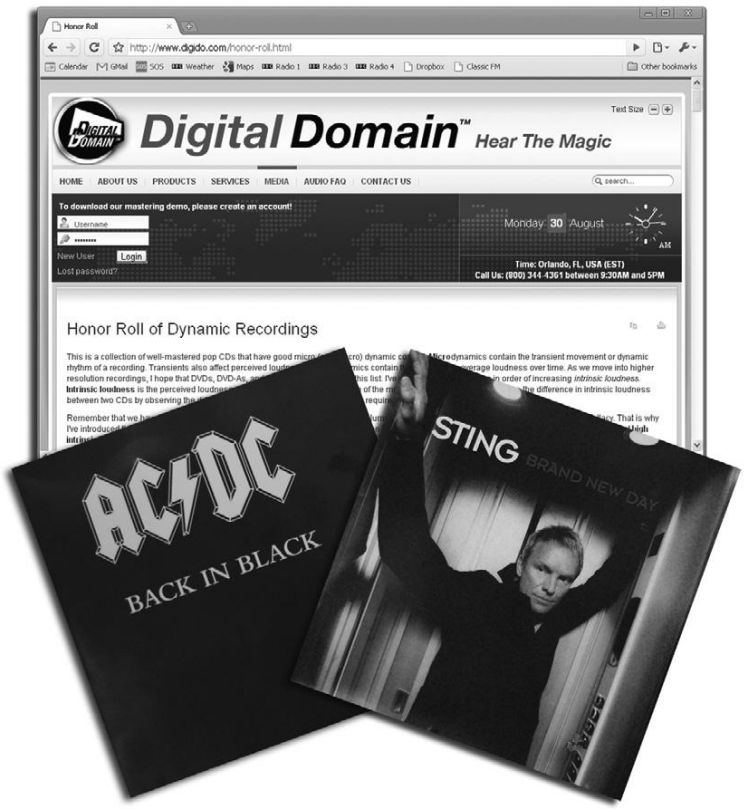

FIGURE 4.5
The Honor Roll of Dynamic Recordings at www.digido.com has a lot of good suggestions for reference material, including albums such as AC/DC's *Back in Black* and Sting's *Brand New Day*.

If I find I've trespassed beyond these extremes, then I'm pretty sure my afternoon Kit-Kat is overdue.

Whether for good or ill, a lot of budget-conscious mix engineers are now taking responsibility for mastering their own productions these days. If you find yourself in this position, then reference tracks can play an important role here too, especially when it comes to the contentious subject of loudness processing. Just the activity of deciding which commercial tracks deliver an appropriate tradeoff between loudness enhancement and its detrimental side effects will be of benefit to you, because in the process you'll make up your own mind where you stand in the ongoing "loudness wars" debate and can choose your own processing methods accordingly. It's even more important with mastering decisions, though, that you make sure you can really hear what you're doing, so don't make any hard-and-fast decisions while whittling down your shortlisted tracks until you've scrutinized them all on a representative range of monitoring systems.

DATA-COMPRESSED FORMATS

Now that so many musicians store their record collections in "virtual" form on iPods, MP3 players, and mobile phones, it's tempting to use these devices for quick-access mix referencing. I'd strongly advise against this, though. The problem is that almost all portable music players use heavy digital data-compression routines to pack the maximum number of music files into the onboard storage space. Although the side effects of this data compression are small enough not to concern the majority of the listening public, they still lower the quality bar significantly enough to compromise the effectiveness of your mix referencing. The bottom line is that using lower-quality versions of reference tracks gives you an easier target to aim for, so you won't work as hard and the competitiveness of your mixes will suffer.

Getting the Best out of Your References

Getting the best out of your mix references isn't just about selecting the right tracks, though. The way you go about making your comparisons is also a vital part of the equation. Because your hearing system is so quick to compensate for the differences between mixes, the quicker you can switch between a reference track and your own work-in-progress, the more revealing the contrast is likely to be. This is one of the times when the source-selection buttons of a mixer's monitoring section or a dedicated monitor controller can be really useful, allowing you to alternate swiftly between your mix and a reference CD. Another good way to sort out the switching is to export your draft mixdown as a stereo audio file and then import it into a new DAW project alongside your reference tracks. But at the end of the day it doesn't matter how you implement instantaneous switching, as long as you make a point of actually doing it.

FIGURE 4.6
Two possible mix-tonality references: Madonna's "Sorry" (from *Confessions on a Dance Floor*) has a lot of top end, whereas Pussy Cat Dolls' "Taking over the World" (from *Doll Domination*) has a lot of low end.

The ear's adaptive capabilities are also the reason why I tend to edit out just the highlights of my reference tracks, because that makes the contrasts between each snippet much more starkly apparent. Even a short five-second intro can heavily skew your perception of how the following full-band entry sounds, for example. A beneficial side effect of this editing is that the files take up less storage space, so they're more portable and many more of them will fit on a single audio CD. If you have to do work across a number of different musical styles, this is no small consideration.

Just as important as instantaneous monitor switching is the ability to loudness-match reference tracks with your mix, as subjective volume differences between the tracks will confuse the evaluation process: it's human nature that anything that sounds louder tends also to sound better, irrespective of whether it actually is. Again, some consoles and dedicated monitor controllers will give you independent control over the level of your reference CD's playback, but you may find loudness-matching easier to do if you've imported your mix and the reference tracks into a new DAW project. That said, there is a big practical difficulty facing small studios in this respect: if you're referencing against the kinds of super-loud modern productions where the side effects of extreme mastering processing significantly impact on the overall mix sonics, it can be almost impossible to judge whether your own unmastered mix is actually delivering the right balance or tone.

In such circumstances, I fall back on two main workarounds. The first technique is to seek out unmastered mixes of relevant commercial releases or else

mastered tracks where the loudness processing has been kept fairly subtle. The former are as rare as hen's teeth as far as most people are concerned, but the latter are slightly easier to find—in Bob Katz's Honor Roll, for instance. The other workaround is to apply loudness processing to your mix while referencing, in order to give you a better idea of, as Kevin Davis puts it, "what the mix sounds like smashed up and brickwalled from mastering."[23]

Dylan Dresdow mentions doing something similar when working with the Black Eyed Peas: "Pseudomastering gives me an idea of what elements potentially suffer during the mastering stage, specifically in regards to loudness… But [I send] everything to the mastering engineer without any processing."[24] (I'll offer some practical pseudomastering strategies when we get to that part of the mixing procedure in Chapter 19.) Many engineers also deliver a loudness-boosted playback version of their mix to clients so that they can more easily compare the mix against their own favorite commercial records. However, it is nonetheless essential that you follow Dresdow's example and retain an unprocessed version of the final mixdown file so that you (or a specialist mastering engineer) can refine or redo that processing at a later date.

BEYOND MIXING: OTHER ADVANTAGES OF REFERENCE MATERIAL

As if improving your mixes weren't a big enough reason for building up a decent selection of reference tracks, the benefits actually reach far beyond mixing. First of all, the simple process of selecting the audio is excellent ear training and helps familiarize you with the sonic templates of different styles. Then there's the fact that you'll inevitably become familiar with how your reference material sounds on a large number of different systems, and this means that you can start to judge new listening environments in relation to this body of experience—a lifesaver if you regularly do serious work in unfamiliar studios or on location. As John Leckie notes, "What I do whenever I go into a strange control room anywhere in the world [is] put my favorite CDs on, and I soon know exactly what the room is doing."[25] Mike Stavrou stresses the importance of prompt action here:

> Use your reference CD before you get used to the coloration of any new room and its monitors. When you first walk into a new control room, this is the moment your ears are most sensitive to the room's acoustic anomalies. After four days you will acclimatize and in so doing will slowly become deaf to the idiosyncrasies and "nasties" of the room.[26]

Trusted test tracks can also sort the wheat from the chaff when you're auditioning monitor speakers, or indeed any other audio playback component. This is where the Bob Katz Honor Roll really scores for me. Many of these albums will reveal small playback distortions much more starkly than the overprocessed masters that infest the mainstream charts.

4.3 EVERYONE NEEDS A SAFETY NET

So far in this book I've described how to set up sensible nearfield and supplementary monitoring systems so you can hear what you're doing, and I've suggested many ways to improve the reliability of your mix decisions. If you can apply these ideas in your own small studio, then you'll have already taken the biggest step toward achieving competitive, professional-sounding mixes 99 percent of the time. However, unless you've sacrificed your eternal soul in exchange for supernatural arrangement and mix processing skills, there's no getting away from the fact that you're human, and humans have off days. No one, but no one, escapes the occasional handbagging from Lady Luck—not even the top names in the industry. This is why it's common practice for the professionals to hedge their bets when delivering mixes to their clients.

The most flexible way you can do this is to make notes of all the equipment routings and mix settings used, so that the mix can be recalled in its totality if changes are requested at a later date. However, this kind of recall has always been a bit of a minefield in analog setups, because no mix ever seems to come back the same second time round, no matter how well-documented the settings. Your favorite buss compressor might have blown a valve in the interim. Your new studio assistant might not line up the tape machine quite the same way the previous guy did. A roadie might have drunkenly widdled down the back of the console. It's the entertainment industry, for heaven's sake—anything could happen!

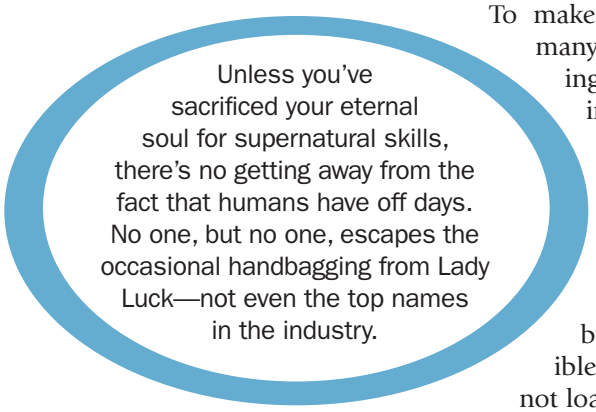

To make mix recall easier, quicker, and more reliable, many mix engineers have now migrated to working entirely "in the box" on a computer DAW mixing system. Although this approach can be much more dependable if you've got serious resources to throw at maintaining the hardware and software, small DAW-based studios often turn out to be no more reliable for recall purposes than their analog counterparts. This is partly because of simple hardware faults and data corruption, but also because of every studio owner's irresistible urge to upgrade—upgrade a plug-in and it might not load into "legacy" projects any longer; upgrade your DAW and the sound of the EQ processing may have been tweaked behind the scenes; upgrade your operating system and you might as well try recreating the mix on a pocket calculator. So while a full mix recall can provide a good solution to mix problems that are discovered in the short term, you'd be loony to bank on that for revisions further down the line. A much more reliable way of managing this kind of longer-term fence sitting is to create alternate versions of your mix—something most pros do as a matter of course. "I'll print as many mixes as needed, depending on how difficult the artist is to please," says Allen Sides. "I just do not want to have to do a mix again."[27]

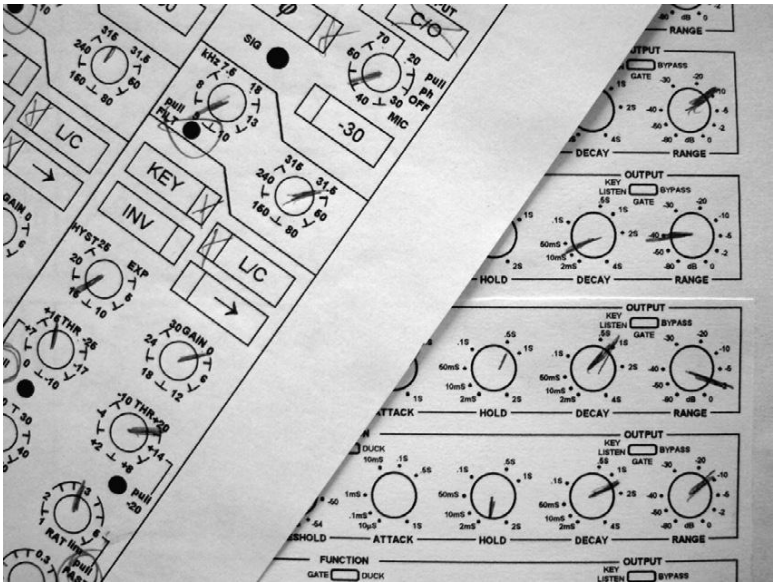

FIGURE 4.7
No matter how carefully you try to save and document your work, mix recall will always be a bit hit and miss—even on computer systems.

Alternate Balances

Once everyone concerned is happy with the general sound of a given mix, revision requests are usually a question of altering the balance, and easily the most common part people want to tweak with hindsight is the lead-vocal level. "It's a hard balance," explains Spike Stent. "Record companies always want the vocal louder than God, but you need to keep the power in the track as well."[28] Because lead vocal levels are so hard to get right, most mix engineers will print a "vocal-up" mix for safety's sake, but there's nothing to say you can't print multiple versions with different vocal levels as extra insurance if you're generally less confident with this aspect of your mix. Bear in mind as well that you could later edit between the different vocal-balance versions if there were just the odd phrase that needed tweaking, so you get a lot of extra flexibility this way.

Bass instruments are also frequently the subject of postmix prevarication, partly because there are often such large differences between the low-end presentation of different playback systems, so alternate versions are commonplace here too. "We'll always do a bass version of some sort," says Chuck Ainlay, "whether it just be a 'bass up' or a 'bass and bass drum up' or just 'bass drum up.' When you get to mastering, sometimes there's an element that EQ doesn't necessarily fix, so they might use a different version instead—like that 'bass drum up' version—to add a little bit of punch."[29] Different bass-balance versions are even more relevant in low-budget setups, because acoustic treatment

can't ever completely even out typical small-studio room-mode problems. If you carefully restrict your low end to just a couple of tracks in the arrangement, and then do a few alternate mix versions with different balances of these tracks, you should have ample ammo to handle hidden low-frequency gremlins that may surface in the future.

But vocals and bass instruments aren't the only tracks that might warrant alternate mix versions. Any track level about which you're undecided in your final mix can benefit from its own little safety net. "I'll just go through and pick some instruments that somebody might have questioned, 'is that loud enough?'" says Jon Gass, who usually ends up doing roughly a dozen versions of each song as a result.[30] Given that an alternate version will only take a few minutes at most to do, I'd normally err on the side of doing more alternates rather than fewer, but it all comes down in the end to how confident you are of your mix decisions versus how much time you have to spend on speculative ass-saving precautions.

Vocal-out and Solo-out Mixes

There's another type of alternate mix that is frequently requested on solidly practical grounds: a "vocal-out" mix, sometimes also known as a "TV mix" or "public appearance (PA) mix." This is a mix with the lead vocal completely removed, and it's designed primarily for use in live performance situations where an artist is required to perform a song without the support of a live backing band—either the band is not available or the restrictions of a television/video shoot preclude the use of the band's live sound.

If you find you've completely screwed up your vocal levels, you can sometimes recombine instrumental and *a capella* mixes to give you a "get out of jail free" card.

However, the vocal-out mix can also serve a number of other useful functions at the eleventh hour, because it allows you the option of completely replacing the vocals—perhaps A&R have only realized the vocal performance is rubbish now that the mix is otherwise complete, or the band has sacked its singer, or current affairs have rendered the lyrics out of date or tasteless. You can also chop sections of the vocal-out mix into the main mix if you need to delete a stray expletive or even generate a full PG-rated version from scratch. "The instrumental comes in handy sometimes for editing out cuss words and things like that," says Jerry Finn.[31] (That said, it's usually better to create any "clean" version of the lyrics at mixdown time, so you can be sure to catch things like backing vocals too. It also often makes more musical sense to reverse, process, or replace offending words, rather than simply muting them.)

If there are points in your track where an important instrumental solo takes over the lead role from the vocal part, there's a good argument in such cases for including a mix version without this too. Some singers are also featured

MASTER-BUSS PROCESSING

Some styles of music rely heavily on obvious master-buss processing (most commonly compression) to achieve their characteristic sound, and this is an area of mix technique that can be difficult for the less experienced engineer to get right. Also, when processing a complete mix, the unique personality of the specific model of processor tends to make more of a mark on the overall production sonics than if you had used the same processor on just a few tracks. If you're concerned that this element of your mix may be letting you down, then recording versions of your mix with and without any master buss processing can be worthwhile. That way you can revisit your buss-processor settings later to apply the benefits of hindsight, or maybe rent some esoteric outboard to see if that might elevate the sound beyond what was possible with standard fare. A mastering engineer may also be able to trump your home-brew processing by dint of more high-end gear and golden ears.

You should consider carefully whether any alternate versions of your mix should include your master-buss processing as well. For example, if you leave a compressor in your master buss while exporting instrumental and *a capella* mixes, you'll find that you can't just recombine them to recreate the final mix, because the two mix passes will result in different gain-change patterns throughout the song. Mix stems will also fail to reconstitute the mix sound correctly if they pass through master-buss processing, for similar reasons. "I'll turn the buss compressor off," confirms Mick Guzauski, "because of course that compressor reacts to everything."[32]

instrumentalists as well, so you might actually drop two whole parts for the TV mix so that they can play and sing live in the way they're used to doing—think Alicia Keys or Jack Johnson, for example.

Instrumentals and *A Capella*s

An instrumental mix, completely without vocals, is another sensible fallback, again because it's not uncommon for A&R to demand that vocal performances and arrangements be rehashed if they feel that the mixed production isn't quite on the money. But there's another good reason to keep an instrumental backup—because it allows the vibe of the track to be used in more low-key contexts, such as in the background of television and radio shows. This kind of usage may not seem very rock and roll, but it's nonetheless an exploitation of the artist's copyrights, which can generate royalty payments, and extra income is something that draws few complaints from most of the jobbing musicians I know.

The inverse of an instrumental version is what is usually called an *a capella*—a solo mix of all the vocals, complete with their mix processing and effects. Indeed, if you find you've completely screwed up your vocal levels, you can sometimes recombine instrumental and *a capella* mixes to give you a "get out of jail free" card. However, the main reason to keep an *a capella* these days is in case anyone wants to nab the vocals for a remix in another style. Again, this might not be what anyone has in mind at the time of the original mixdown, but that's no reason not to plan for it given that a successful club remix could bring in a pile of royalties several years from now.

Mix Stems

Mix engineers have been merrily creating all these kinds of alternate versions for decades, but more recently a further future-proofing tactic has begun to gain currency for music work: mixing to stems. Roughly speaking, this means submixing related sounds in your mix to a limited number of channels and then recording their outputs to separate files alongside your main mix. For a typical band recording, say, you might expect to create stems for drums, percussion, bass, guitars, keyboards, pads, backing vocals, lead vocals, and effects returns. This procedure allows you to reconstitute your final mix by adding them back together. Jimmy Douglass attests that this is becoming standard practice professionally. A&R departments have now learned to expect this additional level of flexibility, so you'll seem like you're behind the times if you can't provide it for them, he maintains.[33]

The most obvious advantage of mixing to stems is that you get more retrospective tweakability than is afforded by simple alternate mix versions, but that's by no means the end of the story. Stems can extend the options available to club remixers, for instance, by giving them more in-depth access to the sounds on the original record. It also gives a lot more editing and layering flexibility to computer games developers if they want to weave the elements of your music around onscreen action. For professional productions, however, it's the possibility that the track might need to be cheaply reworked for surround sound that seems to underlie much of the enthusiasm for mixing to stems at the moment. So if you also see a market for your music in surround formats, then you may want to hop aboard that bandwagon too.

CUT TO THE CHASE

- If you're sensible with monitoring and acoustics in the small studio, you should be able to hear everything you need for making a commercial-quality mix. To use this information effectively, however, you need to learn how to convert your own subjective impressions of what you're hearing into more objective mix decisions.
- Part of the problem with listening objectively is that your ears interact in complex ways with what you're hearing, effectively presenting you with a moving target, so you need to understand and work around this. Handy tricks include switching regularly between monitoring systems, taking frequent breaks, choosing and varying your monitoring levels appropriately, and gathering the opinions of third parties wherever possible.
- By far the most powerful technique available for objectifying your decisions is the process of mix referencing. The more time and effort you invest in carefully selecting high-quality reference materials, the clearer your mixing decisions will become and the more you'll push yourself to improve your engineering chops. Instantaneous switching between your mix and your reference tracks is vital if you want to highlight the differences between them most clearly, and the loudness of reference tracks should also be matched.

If you find you need to apply mastering-style loudness processing to your own mix in order to allow more representative comparisons against heavily processed commercial records, then it's paramount that you remove this processing when generating your final mixdown files, even if you plan to master the track yourself.

- Everyone makes mistakes, and professionals respond to this reality by keeping the most important options open for as long as they can. Although full mix recall is now widely available in principle, the time-honored method of printing alternate mix versions is actually more dependable over the long term. Alternate mixes can provide the mix engineer with greater peace of mind regarding critical balance and processing judgments, and they allow the producer enough flexibility to rearrange or reperform critical lead parts should these require further attention post mixdown. Artists can benefit too, both from additional live performance opportunities and from more widespread exploitation of their copyrights.
- A more recent trend is to print submixed stems from which the mix can be recreated. Mixing to stems can provide more postmixdown tweaking capacity, extra copyright-exploitation possibilities, and a certain amount of future proofing if there's a likelihood that you'll have to put together a surround-sound mix in a hurry or with a minimal budget.

Assignment

- Find a simple way to switch between your available monitoring systems and adjust their levels so that it's easy to make this a habit.
- Create your own reference CD, using only audio sources that use no data compression and evaluating the merits of any potential track on all your monitoring systems.

www.cambridge-mt.com/ms-ch4.htm

This page intentionally left blank

PART 2
Mix Preparation

I know what you're thinking. "Enough of the monitoring stuff already; let's get mixing!" That's a fair comment, but for the benefit of anyone who skipped straight over Part 1, let me be absolutely clear: you've missed out the most important bit of the book! Nothing from here on will help you much unless you can actually hear what you're doing.

Even with that bit of tub thumping out of the way, though, it's not quite time to start flexing the faders yet, because one of the key factors in putting together a professional-sounding mix is preparing the ground properly. A major stumbling block for small-studio owners is that they're usually trying to mix the unmixable—their problems derive at least as much from shoddy mix preparation as from any shortcomings in their processing skills. This aspect of production has rarely been touched on in print, so most amateur engineers underestimate its significance. But don't be fooled: proper mix preparation is fundamental to achieving commercial results in any kind of repeatable way.

A major stumbling block for small-studio owners is that they're usually trying to mix the unmixable—their problems derive as much from shoddy mix preparation as from any shortcomings in their processing skills.

But if mix prep's so critical, why don't members of the industry's top brass talk about it more? The first answer to this perfectly justified question is that the best mixing engineers tend to work with producers

and recording engineers who understand how to track, overdub, arrange, and edit their projects into a state where the mix engineer can concentrate on the creative aspects of mixing without distractions. "My goal is always to give something to a mixer where he can push up the faders and it's reasonably close to the architecture that I want in terms of the arrangement," says Glen Ballard, for example. "If you give them that to start with, it frees them up to really focus on detailing the sounds and making everything sound great. I try not to add to their burden the idea that they've got to arrange the song as well."[1] With the demise of many large studio facilities, fewer small-studio users have now had the benefit of the traditional "tea-boy to tape-op" studio apprenticeship and simply don't realize the amount of time and effort expended on professional productions before mixdown.

Because quality-control procedures are by nature usually quite tedious, it's also understandable that high-profile engineers tend to gloss over them, not least because they often have the luxury of farming out heavy editing and timing/pitch-correction donkey work to specialist DAW operators behind the scenes. And although shrinking recording budgets are increasingly exposing even the biggest names in mix engineering to questionable home-brew production values, it's their assistants who actually bear the brunt of sorting this out, so the boss won't usually dwell on it for long. "The first thing that happens is that Neil Comber, my assistant engineer, will prepare the session for me, cleaning it up and laying it out in the way I like," explains Cenzo Townshend. "All in all, Neil usually has three hours of work to do before I can actually begin the mix."[2] Many other top-tier engineers have also briefly described their own assistants fulfilling broadly similar roles: Michael Brauer,[3] Jaycen Joshua,[4] Chris Lord-Alge,[5] Manny Marroquin,[6] Tony Maserati,[7] and Spike Stent,[8] among others.

Because mix preparation is so important, and a matter of course commercially speaking, Part 2 of this book explores it properly within a small-studio context. My aim is to put you on more of a level footing with the professionals when you actually reach the mixdown itself, because that's almost always half the battle where self-produced records are concerned. The more of this kind of technical work you can get out of the way in advance, the more you'll be able to stay in touch with your vital creative instincts during the mixing process itself.

CHAPTER 5

Essential Groundwork

5.1 STARTING AFRESH

For much of this book I'm discussing mixing as if you're working with other people's musical material, and in such situations it's usually easiest to start with a blank slate as far as mix processing and effects are concerned, rather than working from the client's rough-mix setup—even if you have a mixing system identical to theirs. There's more than just convenience to this, though, as it also allows you to bring some valuable fresh perspective to the production, responding to what's actually important without preconceptions. This is why, even if you're actually recording, producing, and mixing everything yourself, I strongly recommend that you still treat the mixdown stage as a separate task. "I deliberately define the point at which I start mixing," explains Fraser T. Smith. "I think it's helpful for my headspace…. I'll clear the decks, clean up the session, consolidate all the tracks, and look at the session purely from a mix point of view…. It feels like the session is set in stone, and nothing is going to be moved any more. It's a visual/psychological thing."[9]

My favored way of doing this is to prepare a set of unprocessed raw audio files, one for each track in your arrangement, much as you'd do if you were having your music mixed by someone else. Make sure to bounce down the output of any live-running MIDI instruments as audio too. This reduces CPU load (allowing you more power for mix plug-ins), discourages endless tinkering with synth settings during mixing (although you can still rebounce an altered version of the part later if you really need to), and also avoids an insidious problem with some MIDI instruments where they respond slightly differently with every play-through—there are more than enough mind games to deal with at mixdown without this kind of thing going on. "I never run anything live, from a sequencer," affirms Tom Lord-Alge. "I don't want to have to worry about synchronization or issues of sound level."[10] Splitting out the individual sounds from virtual drum instruments and multitimbral samplers is a good idea too, because it gives you maximum mixing flexibility.

If there is some effect that is an essential musical feature of the track (perhaps a pitch corrector, a specific guitar-amp simulator, a synth-style resonant filter,

or some unique special effect), by all means leave that in when generating your bounced multitrack files, but bear in mind that it can save time later if you bounce out a clean version as well, just in case the effect needs remodeling to work properly in the final mix context.

5.2 ENHANCING NAVIGATION

One way or another, you should end up with a fresh DAW project file containing a bunch of bounced audio tracks that all start at the same point. The first step is to make it as easy as you can to get around the place, and this is something many small-studio users seriously undervalue. It might seem petty, but the last thing you want to be doing when you're mixing is waste time working out where the blazes that second tambourine overdub got to. Or resetting the EQ of the wrong guitar part by mistake. Or editing the synth pad out of the wrong chorus. Creative thoughts are fragile little things, so any kind of delay between thinking them and acting on them while mixing is really bad karma. What was I going to do with that second tambourine now? Oh well. Never liked it anyway.

Organizing Your Tracks

One time-honored way of making your navigation of a project instinctive is to standardize the track layout of the instruments in your arrangement so that you instinctively know where to find your kick drum, bass guitar, lead vocal, or whatever when inspiration strikes. Part of this process involves putting the most important instruments where you can easily access them and submixing any large groups of subsidiary parts to fewer mixer channels so that controlling them en masse is easier. Michael Brauer's views are emblematic of most of the biggest names in mix engineering here: "Whatever the session is, it will be a maximum of 44 tracks…. I am not going to mix 200 tracks on the desk for you, so the best thing for you to do is to give me the stereo blends that you like…. I'm happy to mix 16 tracks. The simpler the better. I don't want to have to fight my way through a mix."[11]

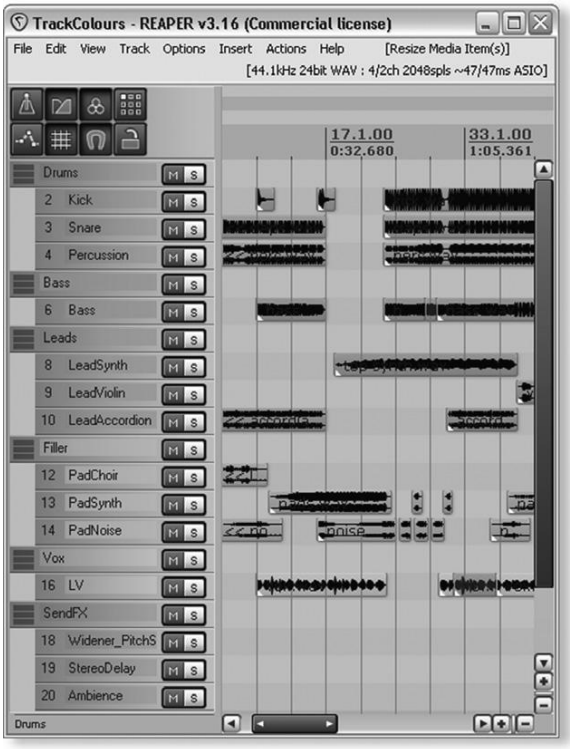

FIGURE 5.1
Laying out your tracks in a logical and consistent manner can seriously speed up your work, as well as freeing up your brain for more creative thinking during the mix.

Many engineers, including both Lord-Alge brothers, also limit their track count in a similar way, with certain sounds always arriving on the same mixer channels and the most important sounds closest to hand. "If everything is parked in the same place," explains Chris, "all you have to worry about is the song. When you're mixing you want to eliminate

all the things that make you think outside of the song."[12] Tom elaborates on his own preferences: "I like to have my drums on tracks 6 to 14. The bass guitar will always be on 15 and 16. Channels 17 to 24 contain the main instruments: guitars, keyboards, or whatever they are. The lead vocal will be on 25 and any additional vocals go after that. Any music that's left over will go on 33 and upward. [Percussion] will go to channels 5 and down. … The faders for channels 17 to 24 are close to my left and 25 to 32 are close to my right. Tambourines can live on channel 1—I don't want to have to move over there all the time; I don't need that much exercise! Of course, by centering the most important instruments on the desk I can also remain in the ideal listening position between the monitors for most of the time."[13]

Naturally, the "in the box" DAW mixing systems typical of many small studios make physical ergonomics less of an issue, but the general principle of standardized track setups still holds: if you can stick fairly closely to some kind of generic layout with each successive mix project, you'll clear headspace for making the really big mix decisions. Practically every mixing program now has sophisticated organization systems for your tracks, so get busy with them to create a visual connection between all your drum tracks, say, or to provide easy global control over massed rhythm guitars or backing vocals.

If you can stick fairly closely to some kind of generic session layout with each successive mix project, you'll clear headspace for making the really big mix decisions.

Colors and Symbols Speak Faster Than Words

As part of organizing your tracks for mixing, you should also get into the habit of labeling your tracks sensibly. If you have 15 tracks just labeled "Guitar," you'll forever be tweaking the acoustic instead of the power chords. More mind games; less hair left. So do yourself a favor and try to make the track names mean something. On many DAW platforms, it's also good practice to abbreviate the names as much as you can so that you can still read them when they're truncated to fit across the width of the channel strip in your software mixer.

However, most software users have an even more useful tool for track identification: coloring. Our brains are much quicker at responding to colors than to words (which is why traffic lights use colors, for example), so the more you can color-code your tracks and audio files, the quicker you'll be able to navigate around them. As Robert Orton has noted, "I'll start my mixes by laying out the session in the way I like and color-coding everything, so it's easier to orientate myself and I instinctively know where things are."[14] If all your drum parts are yellow and all your bass tracks red, then you always know which type of audio you're looking at even when you've zoomed right in to edit individual audio waveforms or when your track names are covered by other floating windows. Jochem van der Saag has another useful

suggestion: "[I] add colours to remind myself that I need to do something, so I may mark a section in a weird colour so next time I open the window I know I have to fix it."[15]

Graphical symbols are usually quicker to comprehend than text too, and they can be understood even when very small on screen, Therefore, if your DAW gives you the option to add track icons of any type, then these might also be a useful visual aid.

Dividing the Timeline

Mixing typically involves a lot of listening to small sections of your mix, as well as a lot of comparing different song sections against each other, so you can save a lot of aggro by having a clear idea of the musical structure of the production and how it relates to the minutes and seconds of your playback timeline. On dedicated hardware recorders, you'd have had a scrap of paper with track times scribbled next to a list of the song sections, but in the software domain you normally have a much more elegant scheme whereby you can display song sections as bars or markers in the project's main horizontal time ruler. Getting this chore out of the way at the first possible opportunity is especially important if you're working on someone else's production, because it significantly speeds up the process of getting to know the music. Again, naming and coloring these section markers will only help you zip about the place more swiftly.

If there's any work to be done in terms of synchronizing your sequencer's metric grid to that of the music, then that's also well worth sorting out before any proper mixing starts. This kind of thing is usually tedious to do, and trying to tackle it right in the middle of the mix is a recipe for inspiration nosedive.

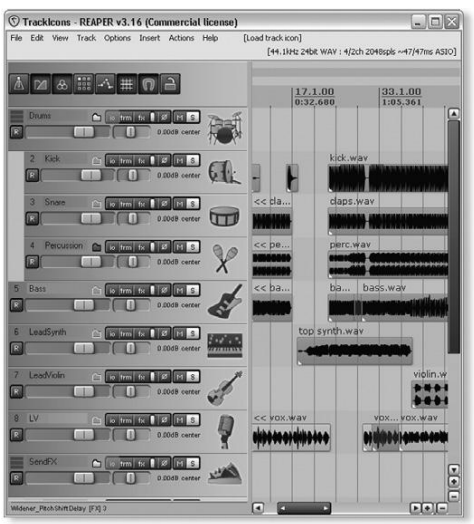

FIGURE 5.2
Track icons can help you identify tracks much more quickly than text names.

5.3 PROJECT RECONNAISSANCE

Spotting Trouble and Hidden Gems

Once you've made sure you can find your way around your mix project as quickly as possible, it's time to start listening through to the tracks individually so you can begin to build a picture of what's there. Even if you recorded the whole production yourself or you're already familiar with a previous rough mix, there's a lot to be gained from listening to at least a snatch of each multitrack file on its own before you properly start work with the mixing. For a start, it's a good opportunity to edit out any silences in your audio files, which not only reduces strain on your computer resources but also typically makes the spread of recorded material across the timeline clearer to see. Plus, you get the opportunity to deal with purely technical issues (such as unwanted background noise, audio glitches, and

lumpy audio edits), which might otherwise slip through the net when everything's going at once. "I will always begin a mix with cleaning up the session," says Demacio Castellon, "doing crossfades, making sure there's no headphone bleed or pops. I hate pops! I spend a lot of time making sure everything is smooth and clean."[16]

Another benefit to be gained from checking through the individual tracks before getting down to mixing is that this process often unearths moments of magic that have been all but hidden by competing arrangement layers. Maybe it's a fantastic little string-bend fill from one of your guitarists, a cheeky little ad-lib tucked away at the end of a backing vocal track, or some choice words from the drummer as he spills coffee over himself—basically anything that makes you laugh, smile, boogie, or punch the air. One of the most surefire ways to win over most clients is to dramatically unveil forgotten gems like these in your final mix, because they can make the mix sound fresher without you actually having to add anything. In fact, several times I've been congratulated for "putting a new part" into a mix, when all I've done is dredge up one of the client's own buried sounds that had long since faded from their memory. The beauty of this trick when mixing other people's work is that it's low-risk, because anything in the original multitrack files implicitly bears the client's seal of approval, whereas there's nothing like that guarantee for any truly new parts you might add.

Again, do make use of your DAW system's coloring tools to highlight these kinds of special moments, because they're usually scattered thinly across a variety of tracks in most productions and you'll lose track of them all otherwise. There are few activities more soul-destroying mid-mixdown than searching laboriously through 24 tracks of vocals for that one killer vocal ad-lib. (It always seems to be on track 23, and after all that it actually turns out to have the wrong lyrics.)

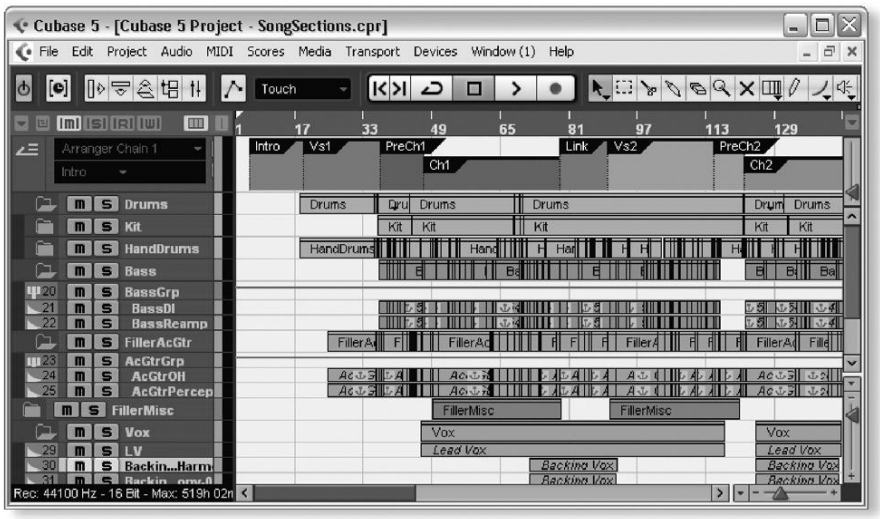

FIGURE 5.3
Use the markers within your own software to name the sections of your song; otherwise it's easy to confuse yourself about which chorus you're actually working on, especially when your view is zoomed a long way in for fine editing or automation work.

Multing

The final thing that you can start doing during your project reconnaissance is any multing you think might be necessary. In the context of a DAW mixing system, the word "multing" primarily refers to the procedure of chopping up a single audio file and distributing the pieces across several different tracks, thereby allowing each file section to benefit from its own mix settings. This simple technique is a terrifically powerful ally of the small-studio engineer, and given the enormous track counts now available on modern computers, you can afford to go on multing till your nose bleeds!

The simplest application is where any instrument plays significantly different parts in different sections of a song. A single acoustic guitar, for example, may well need very different treatments for its verse finger picking and chorus strumming, so multing those sections to different tracks makes great sense. It's also common for important parts such as lead vocals, rhythm loops, or bass instruments to be multed over several tracks so that their processing adjusts to adapt to section-by-section changes in the backing arrangement.

"I use many different EQ and delay settings in many different parts of the song," says Spike Stent, for example. "So when you have one vocal that goes all the way through a track, it may be multitracked three or four times to different EQs and different effects…. You will normally find that the verse and chorus vocal sounds are entirely different."[17] Mike Shipley is willing to push this concept to extremes if necessary: "It's just a matter of making the voice sit in the right place, and if it takes ten channels to do it with different EQ for different parts of the song, then that's what it takes."[18]

Another frequent target for multing in my mixes is backing vocal parts. A single backing vocal track can easily have several phrases within it that all serve different purposes, especially within dense pop or R&B vocal arrangements: verse

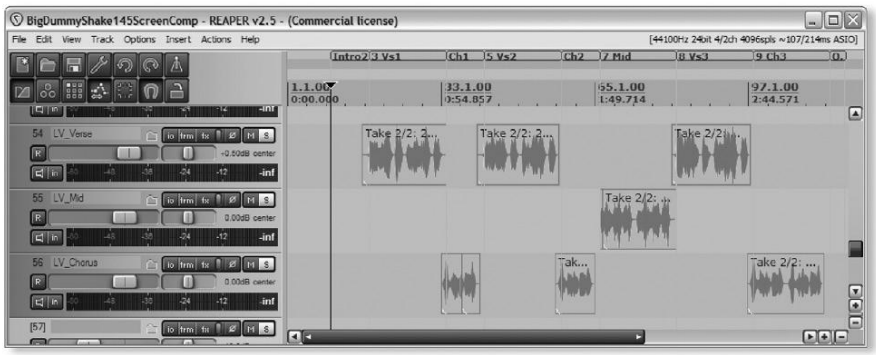

FIGURE 5.4
Multing the lead vocal part, as in this screen grab, is common practice, because it allows you to adapt your mix processing to different sections of the arrangement.

"oohs," prechorus ad-libs, chorus harmonies, whispering in the middle section—the list of possibilities goes on. Multing out all these bits and grouping tracks of similar phrases together is often the only way to make sense of things at the mix.

It's impossible to predict everything you might want to mult until you're actually mixing, so be prepared to come back to it at any later stage, but if you can handle the bulk of obvious mults before the mixdown, it'll not only speed up the process of familiarizing yourself with the material on disc, but it'll also save unnecessary interruptions of your mixing flow later on.

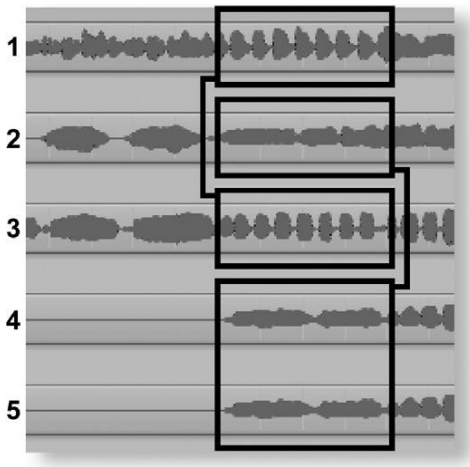

FIGURE 5.5
A section of a backing-vocal arrangement. Notice how the parts on tracks 2 and 3 are similar only to start with—track 2 then matches tracks 4 and 5, whereas track 3 matches track 1. Multing these parts would make sense here to allow similar parts to be processed together.

5.4 CORRECTION TIME!

If you've followed through all the steps in this chapter, locating the track and time location of anything you hear in the mix should be a no-brainer—which is exactly what you want it to be, because you've got much bigger fish to fry. You should also have a smattering of tracks and audio segments that you've colored to indicate that they need special attention, whether that's corrective processing or just some added emphasis in the mix. The next stage is to attend to any timing and tuning concerns, an activity that warrants close scrutiny in the next chapter.

CUT TO THE CHASE

- Small-studio mix engineers frequently overlook proper mix preparation, leading to a frustrating mixing experience and, in some cases, productions that are effectively unmixable. The only way to get on comparable footing with the professionals is to do your groundwork.
- It's usually best to start a mix with a clean slate, working from raw multitrack audio files in a new DAW project, even if you're engineering and producing everything yourself.
- Speed up your navigation of the mix project as much as you can so that you can concentrate more on critical balancing and processing decisions. Standardize your track layout to make finding the appropriate controls more instinctive, and make as much use of track naming and coloring as you can to avoid confusion. Set up timeline markers to indicate the main mix sections and deal with any synchronization issues as soon as you can.
- Listen through to the individual multitrack files to remove silent regions, fix purely technical problems, and highlight any audio sections that need timing/tuning correction or special emphasis in the mix. While you're at it, mult out any sections of tracks that you suspect might benefit from independent mix settings, but don't be afraid to carry out further multing later on during the mix process if required.

Assignment

- Find a mix that you can work on while progressing through the rest of this book, and create a fresh DAW project file for it.
- Use colors, icons, markers, and text labels to speed up project navigation and to highlight tracks or audio sections that need special attention.
- Use multing wherever appropriate.

www.cambridge-mt.com/ms-ch5.htm

CHAPTER 6

Timing and Tuning Adjustments

Everyone knows that a great live performance has the capability to connect with a listener's emotions, in spite of (or even because of) numerous tuning and timing discrepancies. However, for mainstream productions designed to withstand repeated listening, commercial expectations are now extremely high regarding the tightness of both pitching and groove, and much time and energy is routinely expended buffing these aspects of many professional recordings. Everybody's heard songs on the radio where this tinkering feels like it's been taken way too far, but opinions clearly differ widely on what "too far" is. At the end of the day, how much corrective processing you use in your own work remains a personal judgment call. However, I'd urge you to resist any dogmatic ethical stance on the subject, and judge the merits of any such studio jiggery-pokery purely in terms of whether it serves the music. If an ostensibly out-of-tune note encapsulates the emotion of the music, then have the courage to let it lie. If the most emotional vocal take has some distractingly duff pitching, then consider that the general public (and indeed the artist) will probably thank you for tastefully drawing a veil over it.

If an out-of-tune note encapsulates the emotion of the music, then let it lie. If the most emotional vocal take has some distractingly duff pitching, then the general public (and indeed the artist) will probably thank you for tastefully drawing a veil over it.

Though the audio fidelity of pitch/time-correction processors has improved tremendously during recent years, none of the current crop of algorithms can claim to be completely free from undesirable side effects, and some pundits use that as an argument for rejecting their use out of hand. However, any mix processing can display undesirable side effects, and it's simply the engineer's responsibility to evaluate whether the cure is better than the disease in each specific instance. The slightest pitch-correction artifact on a classical solo piano recording may not be appropriate, whereas even some appalling processing nasties may be a minor concern when tuning one of many wayward background vocals within a busy pop mix.

Mixing Secrets for the Small Studio.

That said, having heard literally hundreds of amateur mixes, it's my firm opinion that most users of small studios don't pay nearly enough attention to timing and tuning touchups, even taking into account the different demands of different styles in this department. This is doubtless partly because these tasks are about as invigorating as watching paint dry, and I can sympathize wholeheartedly there, but it's one of those bullets you've got to bite if you want to compete commercially these days. Those inexperienced with corrective editing also tend to stop short of what's required because they're unable to achieve enough correction without incurring excessively unmusical side effects, and that's the situation I want to try to remedy in this chapter. With a bit of care it's perfectly possible to achieve adequate timing and tuning correction in practically any mix without giving the game away by making things sound unnatural. Certainly, I use a degree of corrective editing on almost every mix I do, but I've never had any complaint about it—in fact, I can count on the fingers of one hand the occasions anyone's even commented on it at all.

6.1 GROOVE AND TIMING

There's a persistent studio myth that tinkering with performance timing will inevitably sap the life out of a production. What this doesn't acknowledge is that there's a big difference between a great natural-sounding groove and a simply sloppy performance. Your drummer can be directly channeling the soul of James Brown, but if the rest of your instruments display the rhythmic panache of Mr. Bean, then you'll still empty the dance floor pretty quickly! True, if you try to turn the drummer into a drum machine, then you'll probably lose more in vibe than you gain in consistency; but if you edit the other parts such that they agree better with the drummer, then the sense of natural rhythmic flow will usually shine through much more strongly.

My top tip for making the best of a production's groove is to work out which single instrument embodies it best, solo it to smooth out any stumbles or areas of inconsistency in that instrumental part if necessary, and then use it as a reference point as you reintroduce the remaining instruments one at a time, tightening the timing of each as you go to taste. In a full mix, it can be all but impossible to work out which combination of instruments is causing some perceived timing unevenness, but if you strip back your arrangement and then progressively rebuild it like this, you can be pretty certain at any stage where to point the finger of blame whenever the groove falters.

For most modern productions, it's usually a no-brainer to start the editing process from the drums, especially if these are MIDI parts or loops, which by their nature flawlessly maintain a consistent rhythmic momentum. But even with live bands, the drummer is typically the player with the best sense of rhythm, so it stands to reason that the drummer's parts will probably offer the best guide. In arrangements without drums, any other rhythm-section part can operate as the main reference point, but it pays in such cases to audition each contender on its own to properly weigh up its appeal before making up your mind. Whatever

part you decide to latch on to, though, you should naturally ensure that the quality and consistency of its groove are as good as possible before you start lining anything else up to it. This is the part of the timing-correction process where there's the greatest risk of oversanitizing things, but fortunately there are a few guidelines that help ward off this undesirable outcome.

Timing: A Relative Perception

The most important thing to remember about the perception of timing and groove is that it is entirely relative. You judge the timing of every new rhythmic event based on expectations you've built up from the patterns of events preceding it—not based on how it lines up with any notional metric grid you might have superimposed on those events. For example, even if a drummer records his parts along to a click track (and therefore ostensibly in sync with your DAW's metric grid), it's perfectly possible for a drum hit that looks as if it's been played late according to your DAW's bars/beats ruler to actually sound as if it's been played early. This might sound paradoxical, but let me explain.

Let's say you've got a song at 100bpm, and the drummer hits the downbeat of bar 1 perfectly in time with the DAW's grid but then doesn't quite keep up with the click track's tempo—for the sake of example we'll say he sets off at 98bpm. If you're listening to the drummer without the click, the groove of his drumming will sound perfectly fluid as long as he remains true to his 98bpm tempo. However, if toward the end of bar 2 the player notices that he's now roughly a 32nd note late compared to the click track he's hearing in his headphones, he may begin correcting himself, advancing bar 3's downbeat hit slightly so that it's only a 64th note late compared to the DAW's metric grid. However, although it still looks late in the DAW, what you'll hear if you're listening to the drummer without the click track is that the downbeat hit of bar 3 appears to be rushed compared to the consistent 98bpm groove you've heard so far. (If you're still scratching your head about this, then check out Figure 6.1 for further clarification.)

The clearest lesson to be learned from this is that you can't afford to judge issues of groove just by looking at your sequencer's metric markings; those floppy things on the sides of your head must always claim the final veto. But there are other ramifications. First, you need to understand that you can't properly evaluate the timing of a given bar or beat unless you audition from a point well before it—you'll usually need to allow a run-up of at least two bars for your own internal musical pulse to latch firmly onto the track's groove. Otherwise you won't pick up those small lurching sensations that alert you that there's work to be done. You also have to realize that when you hear a note as being out of time, the best remedy can actually often be to edit the notes before it. Once you've carried out the edits, you may also have to adjust subsequent notes too, because your earlier edits may have substantially altered the rhythmic expectations of the listener.

It's particularly easy to get yourself into difficulties if your key rhythmic instrument part has been edited together from several different takes or constructed

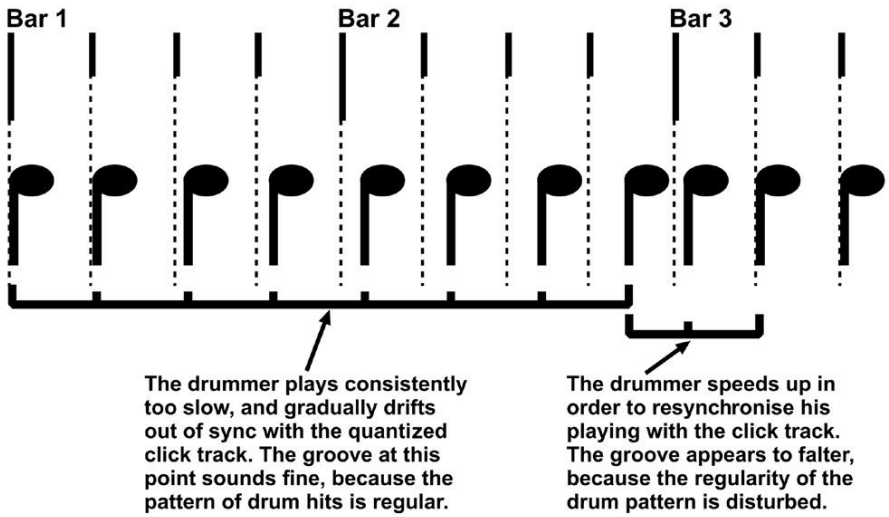

FIGURE 6.1

The importance of relative timing to groove.

from a few repeated performance sections or loop samples. Even if all the sections in isolation are groovier than a purple velvet suit, you can often find that the first couple of hits after each section boundary sound out of time, because each section has a slightly different relation to the main metric grid—in other words, some are slightly behind the beat and others slightly ahead of the beat. Most people's first instinct in this scenario tends to be to edit the hell out of the apparently out-of-time individual hits, but this rarely yields adequate improvement and usually compromises the overall grooviness into the bargain. A better solution, of course, is to slide the timing of each section as a whole while listening for rhythmic lumpiness across their transitions.

Once you understand and apply these underlying concepts, you should be equipped to get on with timing edits without falling for too many psychological mirages, at which point the quality of your work will largely depend on how much practice you get. The more you do corrective work, the more attuned you become to the subtle stumbling feelings that indicate the presence of an out-of-time note. Fortunately (or not, depending on how you look at it), there is usually no shortage of material to practice on in most low-budget productions, so you should get up to speed fairly quickly. It's worth noting that it can take a little while for you to get yourself sensitized to small timing variations at the start of an editing session, so it's not a bad idea to double-check the first few sections you work on after you've had a chance to warm up.

Tightening the Timing

Once your key instrument is fighting fit and ready to go, you can start to rebuild the rest of the arrangement around it, and I suggest starting this process

with those parts that have the most important rhythmic roles. If you've used the drum parts as your initial timing reference, then the bass is often the best part to deal with first, before moving on to any rhythm-section guitars or keyboards. With those parts tightened up a bit, you'll often find that the impetus of the groove is strengthened enough that the remaining tracks don't actually need to be tweaked much at all and may actually help retain more of a natural-sounding vibe.

It's not uncommon in a lot of mainstream productions to layer multiple recordings of the same part in order to create a thicker and more homogenized texture. Rhythm-guitar and backing-vocal parts are most commonly subjected to this treatment. The problem with these kinds of parts from the perspective of timing correction is that once the first layer is edited and in place, it's more difficult to hear timing discrepancies in the second layer. My advice here is to tighten the timing of each layer one at a time, without listening to them together. Once each layer seems to fit the groove in isolation, then you can combine them all. I find that there are usually then only a handful of minor adjustments still to be made and the whole job's done in half the time.

How much of each part you tighten up and how precise you are about matching it to the groove are also key aesthetic decisions, although I have to say that I prefer to err on the side of more accurate timing where rhythm-section parts

WHEN NOTHING GROOVES

For most of this chapter I'm assuming that something in your arrangement has a fairly consistent groove, and for most small-studio arrangements this isn't a bad assumption, given how many tracks these days are constructed from MIDI files or packaged sample loops. However, small-studio band recordings are more than usually prone to the curse of the talentless drummer, and such performances can be difficult to improve with normal corrective edits because the timing is so inconsistent. You can't maintain a groove that isn't there in the first place!

So what to do? Well, probably the least time-consuming tactic is to ditch all but the most passable sections, loop-audition each one to spot-fix as many internal problems as you can, and then cobble together a final take with copy and paste. You'll sacrifice some variety in the player's drum patterns, but at least you'll salvage some semblance of rhythmic momentum. You can sometimes get away with this dodge even if the drummer didn't play to a click.

If there's ostensibly a metric grid around which the drummer flailed hopelessly, then another last-ditch option is to slice the performance into individual hits and quantize them. In a situation where there's no groove to kill, making your drummer sound like a machine can start to seem a lesser evil. However, I'd urge small-studio operators to do their best to avoid this course of action if they value their sanity. First, it'll take you all night; second, a bad drummer's performance will still sound unmusical, even if it's in time; and third, a lot of ropey drummers have a lousily balanced kit, so it'll be a nightmare to mix into the bargain.

are concerned. I've rarely found this to compromise the musicality as long as the original rhythmic reference instrument has a good feel to it. If your reference instrument happens to be a MIDI-programmed part, then extra vigilance may also be called for. "When live elements are being mixed with machine elements, you have to make sure everything is very synced up," recommends Serban Ghenea, "otherwise things like phase incoherence between live and machine tracks can drastically change the sound when you put them together in the mix…. It doesn't defeat the purpose [of having live tracks], though…. You can still hear the little imperfections that make it human."[1]

Many small-studio engineers who produce otherwise commendably tight backing tracks frequently suffer from a blind spot where the timing of lead vocals is concerned. This oversight can really interfere with the way a production swings as a whole.

Extra precision is also usually a virtue with any double-tracked parts that you intend to place on opposite sides of the stereo field, because the brain is extremely sensitive to timing disparities between similar sounds arriving from different directions—even slight flamming between notes/strums can get distracting. Furthermore, many small-studio engineers who produce otherwise commendably tight backing tracks frequently suffer from a blind spot where the timing of lead vocals is concerned. Because the main singer takes up so much mix real estate in a lot of styles of music, this oversight can really interfere with the way a production swings as a whole, so don't spare the editing scalpel where it's warranted. Don't forget to pay attention to where notes end, either. The tail ends of bass notes in particular can really undermine the momentum of the groove if badly timed, and the same is true for other important rhythmic parts and anything else that is up at the front of the mix.

While your DAW's metric grid is a pretty unreliable guide if you're trying to refine the groove of your main reference instrument, that instrument's own waveform profile can actually provide a good visual template for editing the other arrangement layers, so you can save some time at the editing stage if you keep the waveform display of the reference instrument close by as you work. I hasten to add, however, that while carrying out preliminary editing "by eye" like this can significantly speed up your work rate (particularly with layered parts), the process of ratifying your editing choices by ear remains utterly indispensable. You should be fully prepared to shift timing away from what looks neat if your ears smell a rat. Er, you know what I mean!

6.2 AUDIO EDITING TECHNIQUES FOR TIMING ADJUSTMENT

To my mind, a significant reason why a lot of small-studio operators stop short of carrying out all the necessary timing adjustments is that their audio-editing chops aren't up to the task. The software isn't really a limitation in this respect,

as even extremely affordable DAW applications such as Cockos Reaper now have more than enough editing tools to pull off practically any real-world timing edit. The trick to achieving good results is to know where best to edit, what tools to use, and how to hide any editing anomalies you can't avoid.

On the simplest level, performing a timing edit merely involves cutting out the section of audio containing an offending note and sliding it to an improved position, then subsequently overlapping the audio snippets at each edit point and concealing the transition with a crossfade. A crossfade involves fading down the outgoing audio while simultaneously fading up the incoming audio, which avoids the audible click that can otherwise arise if there happens to be an abrupt step in the signal waveform across the edit point.

Camouflaging Your Edit Points

At the risk of illuminating the blindingly obvious for a moment, the easiest way to keep your edit points from being audible is to place the crossfades in gaps in the audio. Although this may seem a thunderingly straightforward observation, what a surprising number of people miss is that small periods of silence are frequently tucked into otherwise seemingly continuous audio, not least those small gaps that typically precede vocal consonants such as "b," "d," "g," "k," "p," and "t." The main editing consideration here is to judge the position and length of each crossfade, but there's rarely any mystery to that in practice.

Obviously noisy sections in an audio signal are equally good spots for edits because the inherently random nature of noise signals tends to obscure any unnatural waveform transitions at the edit point. In these kinds of editing locations you could often use no crossfade at all without any click becoming audible, but I usually put in at least a 2ms crossfade for safety's sake. There are noisy elements in most musical sounds, so there's often a lot of leeway here. For example, vocal parts have breaths and consonants such as "f," "s," "sh," and "ch," acoustic guitar parts have fret slides, and many overdriven electric guitars are so noisy by nature that you can regularly get away with editing them almost anywhere.

When you have to site an edit point in continuous audio, then the best thing is to try to put your crossfades immediately before percussive attack onsets. It might seem like time travel, but even though the transient only happens after the edit, the later event still effectively conceals (or "masks") the earlier one—a freaky psychological audio effect referred to as backward temporal masking or premasking. Here short crossfades of around 5ms tend to work best. Instruments with a clear percussive attack (acoustic and electric piano, tuned percussion, harp, acoustic guitar) will respond best to these kinds of edits, but even the less well-defined attack characteristics of

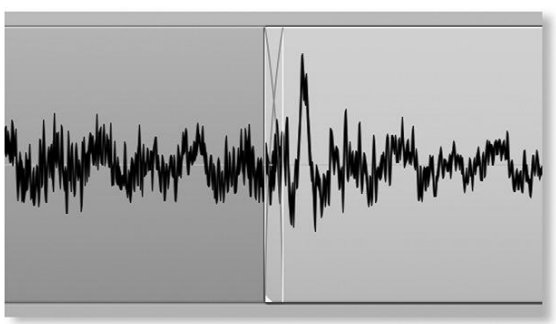

FIGURE 6.2
A premasked edit, just before a drum hit.

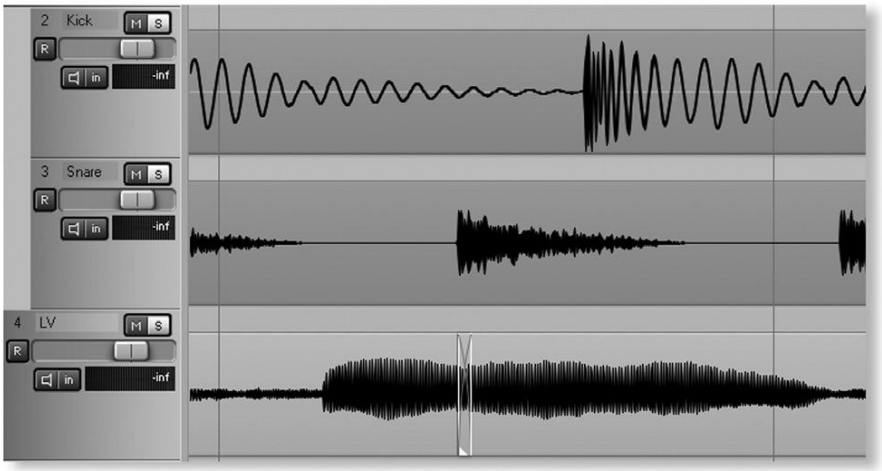

FIGURE 6.3
A masked edit. The snare track can be used to mask even a clumsy edit on the lead vocal track.

big-band and orchestral instruments can be used in this way with a reasonable degree of success.

These three basic types of edit (in silence, in noise, and premasked) are already fairly widely used among the small-studio community, and they'll usually handle 80 to 90 percent of the timing edits you need to do. However, you'll not get 100 percent of the way unless you're aware of two other lesser-known editing possibilities. The first is to place an edit point on one instrument such that a powerful transient or noise signal from another instrument masks it, and this hugely expands your options. Take a stereotypical pop/rock recording, for example. Almost anything could be edited underneath snare hits, lower-range instruments could be sliced under the cover of every bass drum, and longer crossfades on higher-range instruments could hide behind cymbals and open hi-hats. When it hits the chorus, the combined masking effect of drums, cymbals, and high-gain electric guitars is likely to be so powerful that you can sneak a 50ms crossfade edit into any background part wherever takes your fancy.

The second power-user edit really comes into its own for time correction of lead vocals and melodic instrument solos, because it gives you the facility to place well-nigh inaudible edits right in the middle of long held notes. The technique relies on the fact that the waveform of any sustained pitch by its very nature repeats regularly. By ensuring that this repeating waveform pattern continues smoothly across the edit point, you can usually get an edit that is somewhere in the range from serviceable to remarkable just by crossfading over a couple of waveform cycles. It has to be said, though, that there is something of an art in selecting the right location for this type of edit. Notes with vibrato can be tricky to deal with, for example, because the pitch variations alter the length of the waveform repetitions. An edit in the middle of a note that is changing tone over time can also stick out like a sore thumb, and this particularly applies to vowel transitions such as diphthongs in vocal parts—out of

choice I prefer to put matched-waveform edits in "m" and "n" sounds because their dull closed-mouth tone doesn't vary as widely.

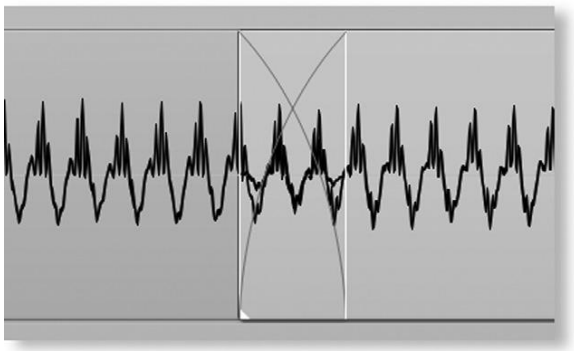

FIGURE 6.4
A matched-waveform edit: essential for exposed edits in sustained pitched notes.

The Role of Time-Stretching

There's one final editing conundrum to be grappled with: whenever you move a note away from its originally recorded position, you create a situation where you have a gap to fill—before or after the edited note, depending on which way you've shifted it along the timeline. If there are already gaps between notes, then one solution is merely to leave a section of silence at the edit point, but this momentary blackout will often sound odd on real-world recordings that have picked up significant levels of background noise, trailing room reverberation, or headphone monitoring spill. In the majority of cases you can just move the edit crossfade somewhere into the body of the note that is too short and then use either a masked edit, a matched-waveform edit, or just a longer crossfade (upward of 100ms) to camouflage the join.

Sometimes, though, none of these options work out: the arrangement doesn't provide any opportunity for a masked edit, the pitch or tone are too variable for a matched-waveform edit, and the track is too exposed in the mix to justify a longer crossfade's phasey-sounding side effects. This is where dedicated time-stretch processing can save the day by extending the note in question, and the good news is that most DAW systems now have built-in processing that is more than adequate for the purposes of gap-filling fixes.

If you leave a section of silence at the edit point, this momentary blackout will often sound odd on real-world recordings with significant levels of background noise, room reverberation, or headphone monitoring spill.

But if your DAW's got time-stretching built in, why bother with editing at all? The first reason is that even the most state-of-the-art commercially available time-stretching technology I've heard still emits unwanted digital yodeling of one sort or another when it tries to deal with polyphonic audio, transients, or noisy unpitched elements of musical sounds, and as you head further down the food-chain these side effects quickly become worse than the timing problems! However, assuming for the moment that you invest in third-party time-stretching software with acceptably benign processing artifacts, I've found that it's actually slower to work with in a lot of cases, for the following reasons:

- You typically have to isolate polyphonic/noisy sections to avoid excessive processing artifacts, and if any artifacts are unavoidable, then you'll have to fall back on normal audio editing anyway.

- It's often much less easy to line up the waveform in an external editor's display with that of your main rhythmic reference instrument, so you can't readily use a visual guide to your advantage in speeding up the editing process.
- The keyboard shortcut support may be less comprehensive in a third-party application than within your DAW, so there's more laborious mousing around to do.

CHOOSING CROSSFADE TYPES

Some DAW systems give you the option of choosing between different crossfade types, which is useful because this choice can make a big difference. If you select the wrong crossfade type, you'll find that the edited track's overall signal level will increase or decrease unnaturally toward the midpoint of the crossfade. As a general guideline, an equal-power crossfade will usually work best for the majority of edits, whereas an equal-gain crossfade will typically give better results for matched-waveform edits. However, if your particular editing software doesn't label its crossfades in this way, don't lose any sleep over it—just experiment with the different types you have and listen for any level unevenness across the edit point. You'll soon get to know which settings provide the smoothest results.

Finishing Touches

Once you've finished any detailed timing editing for each part of your arrangement, it's an excellent habit to quickly run the song down from the top again before adding in the next part. It's easy to lose perspective when you're doing lots of fine edits, and this simple gambit gives you the opportunity to refine more subtle medium-term timing drifts that won't be as apparent when you've got individual notes under the microscope. You may also decide that the overall timing offsets between the tracks would benefit from some adjustment, because the exact relationship between the different tracks will affect the feel of the groove. "An instrument that's playing right on the bar line will feel faster," comments Jack Joseph Puig. "So you find the best place by moving instruments around. I'm talking here about a few milliseconds forward or backward. So within a bar you can have brisk eighth notes or lazy eighth notes, and this is a very important tool when mixing. I always try to find the musicians that are playing more slowly and that are playing more briskly, and I move them around until I get the feel that I want."[2]

It has to be said, however, that overall timing decisions like these do also depend to an extent on the balance between the tracks within the environment of the final mix, so while it's worthwhile adjusting overall track alignments like this at the mix-prep stage, you should be prepared to refine these decisions further down the line if you feel the groove isn't flowing well enough. Mike Stavrou has an interesting take on this, from a psychological perspective: "If a player is rushing, you can slow them down by pushing the fader

up. Conversely, if a player is slightly behind, you can speed them up by lowering the fader."[3] A similar phenomenon is that the perceived overall tempo of a track can be affected by the levels of the less important beat subdivisions within the groove: fade them up and things drag; fade them down and the tempo often appears to increase.

6.3 TUNING ADJUSTMENT

It is a depressing fact that the sour tuning that blights a lot of low-budget productions is often easily avoidable. An electronic tuner can help a lot here, as long as its internal pitch reference is set correctly by ear; digital MIDI instruments can be completely dependable in terms of pitch, as long as their internal pitch adjustments aren't out of whack; and even a pretty hopeless singer's tuning can be improved if you get the headphone mix right. As far as I can see, though, decent tuning (especially on vocal parts) is increasingly being ignored in small-studio productions, and this puts the onus on the mix engineer to fix things by default, because arriving at a competitive mixdown is otherwise pretty much a pipe dream. "If you don't do it, your records will sound strange," comments Steve Lipson. "People's ears have become used to hearing voices perfectly in tune."[4]

Now no one is saying that everything has to be absolutely robotic, but it's nonetheless important to comprehend that lax tuning has important mix implications inasmuch as it affects how well your tracks blend. The more the pitching of your arrangement agrees with itself, the easier it is to make everything seem as if it belongs together. If any track's tuning is out of kilter with others in the arrangement, it will stick out uncomfortably and be difficult to balance correctly, even if things aren't far enough off beam to induce lemon-sucking facial expressions. It's for this reason that so many producers talk about tuning drum kits to the key of the track: if the strongest pitched overtones of the drums fit with the track, you

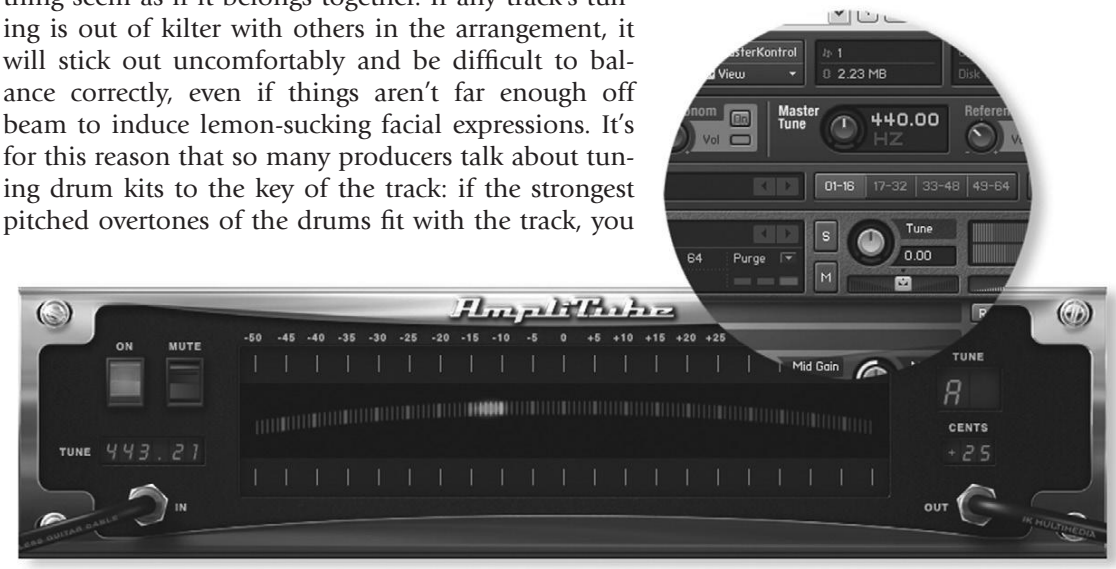

FIGURE 6.5
It's much easier to head off tuning problems while tracking—use a tuner (such as the IK Multimedia Amplitube one shown) for live instruments, and check the Master Tune parameter of any MIDI synths for the best blend.

can fade the kit up further in the mix before it has trouble blending. Tuning is not just a musical issue; it's an important mixing issue too.

However, although unmixably wayward tuning is the most common ailment in low-budget environments, there's also a significant minority of small-studio practitioners who over-egg the pudding by insensitive use of available pitch-correction tools, and that pratfall is just as capable of hobbling the mix by squeezing all the musical life out of it. "Something perfectly in time, something perfectly in tune, could be perfectly boring," warns Steve Churchyard.[5] Steering a course between these two extremes is the art of tuning correction, and the remainder of this chapter offers advice for maintaining this desirable trajectory.

To begin, there is little excuse these days for inappropriate side effects of pitch correction to be audible in a final mix. Again, the built-in pitch processing within current-generation DAW systems is good enough for a large proportion of applications, and excellent third-party software is comparatively affordable for those who feel their work requires more advanced tools or increased processing fidelity. It's not the tools that are the problem—you've just got to understand how get the best out of them.

Choosing Your Targets

The first thing you need to know is what pitch-processing algorithms find easy to handle, namely clean sustained monophonic lines. If there's any noise, distortion, breathiness, or rasp to the pitched-note timbre, this will introduce unpredictable waveform elements that can confuse the pitch-tracking code within the algorithm and cause it to misfire in an unmusical way. Note-attack transients or audio crackles also tend to defy the logic of many pitch-adjustor plug-ins, for similar unpredictability reasons. As for polyphonic audio, until recently your tuning options were limited to fairly small global pitch changes in return for moderately nasty artifacts, even when using top-tier correction products; and even now that Celemony's Melodyne Editor has turned the dream of truly polyphonic pitch-adjustment into reality, its processing still exhibits unpalatable side effects in real-world use, especially if you are unwilling to put in a great deal of setup time to ensure the best results.

If there's any noise, distortion, breathiness, or rasp to the pitched-note timbre, this will introduce unpredictable waveform elements that can confuse the pitch-tracking code, causing it to misfire in an unmusical way.

So in the face of an out-of-tune production, you should think carefully about which tracks to target for correction. For example, even if you only have a couple of out-of-tune guitar overdrive parts, you may well get a better overall sound if you tune the monophonic bass and vocal parts to match them, rather than subjecting these noisy polyphonic sounds to the nasty swirling protestations of overworked digital correction code. Alternatively, if a guitar DI signal is still available, then you might be wise to tune that up instead, subsequently

replacing the errant overdriven parts using a digital guitar-amp simulator. By the same token, it's a great deal easier to replace most problematic MIDI synth parts from their source MIDI files, rather than processing bounced audio for pitch reasons. If the keyboard parts were played live, that still doesn't mean you can't replace them with a MIDI-driven part of your own if that's what it takes to reach a commercial-sounding final mix, particularly as many keyboard parts in mainstream productions are often quite simple.

Once you've decided which parts to process, however, you still need to keep your wits about you. Given that pitch correctors don't like unpredictable waveform shapes, it stands to reason that you should do your best to keep noiselike signals out of their way. Mult the guitar's fret-slides to a separate track. Bypass the pitch corrector for the flautist's breaths. Leave the "s" untouched and concentrate your shifting on the "inging in the rain" bit.

The Right Tool for the Job

You can also improve the quality of your correction work by choosing the right tools for the job. On the one hand, you want to achieve the best sound quality, which in most cases translates to finding the processor with the least distasteful side effects. On the other hand, you want to preserve as much of the original musicality as possible, resisting any alluring machine-automated correction facilities where these risk throwing the baby out with the bathwater.

Although real-time pitch-correction plug-ins are now bundled with many DAW systems, and dedicated third-party software from companies like Antares, Celemony, and Waves is so widely available, there's a lot still to be said for dealing with pitch adjustments using your sequencer's offline editing tools. One definite plus side is that you're probably in less danger of destroying the musicality of a part that way. The main thing that makes a track sound out of tune is if the perceived pitch center of each note is out of line with the context of the track at that moment. As long as the pitch centers feel right, any shorter-term pitch variations mostly just contribute to a sense of natural phrasing and overall humanity. This is why applying simple pitch offsets to offending notes offline can often sound very transparent.

In addition, offline work encourages you to take more moment-by-moment control over the pitch-correction algorithm itself, adjusting parameter settings to suit each shifted note. For example, although most of the time you want to maintain the length of a pitch-shifted section so that the performance timing doesn't change, it's actually much harder on the digital signal processing (DSP) code, so you may encounter unacceptable audio glitches on occasion. If switching off the time correction stops the glitches, you may be able to compensate for the timing change via the time-correction editing techniques mentioned earlier in this chapter and thereby achieve a more successful outcome.

Pitch-correction processors also often apply additional processing to safeguard characteristic instrumental and vocal resonance peaks (called formants), saving

FIGURE 6.6
It's tempting to use real-time pitch correctors like these as a set-and-forget process, but this rarely produces musical or transparent-sounding results.

your singer from transforming into a chipmunk when you're shifting the pitch upward. For small tuning-correction shifts, though, switching off the formant adjustment won't appreciably affect a sound's character in a lot of cases, and it may again improve the smoothness of the corrected sound.

Different offline pitch shifters do vary a great deal in design, though, and many of their parameters are frankly incomprehensible too, so the main thing is just to use your ears while editing. If you're disappointed with the way the processed sound emerges from its transformation, then take the time to tinker with at least some of the available settings. No pitch processing is perfectly transparent, but that's no reason to let more side effects through the net than absolutely necessary.

Automated and Prescanning Pitch Correctors

There is a whole family of pitch-correction utilities that operate in real time, analyzing and correcting your audio as it happens. Antares was the first to really bring this technology to the masses with its original pioneering Auto-Tune plug-in, but since then many other manufacturers have provided variations on this functionality, and they are now built into many DAW systems, such as Apple Logic (Pitch Corrector), Steinberg Cubase (Pitch Correct), and Cockos Reaper (ReaTune). Although these allow you to scrub up a track's tuning in short order using a comparatively small control set, I'd personally advise against falling for the kind of "set and forget" approach they implicitly promote for any important parts in your arrangement. One of the problems

is that they tend to process everything passing through them to some extent, so artifact levels on noisy or transient signals can be more of a problem, and expressive nuances such as vibrato are reined in along with pitch-center deviations by the time you've achieved the desired tuning accuracy. However, although it's possible to improve things to a degree with some parameter adjustment, the fundamental problem is that pitch-detection and correction still require musical intelligence, so the only way to make these processors work really well is to manually assist them in achieving a better result.

"With Auto-Tune you really have to work hard to make it do its thing without being heard," explains Steve Bush. "People seem to think it's some kind of quick fix, but it's not…. For a start it's not always good at detecting what the correct pitch is, but also even if a whole line is consistently a little out, Auto-Tune will always shift note by note."[6]

The most effective way to improve the musicality of pitch correctors is by taking the guesswork out of the pitch detection, because this is the element of the automatic processing algorithm that most commonly errs. There's usually some way to do this with most plug-ins, whether by using MIDI notes or the DAW's mixer automation system to manually inform the plug-in which pitches are arriving at its inputs. Activating the pitch processing only when needed is a good idea too, as there's no sense in flattening out performance inflections where there's no tuning benefit at all. Beyond these measures, the main thing

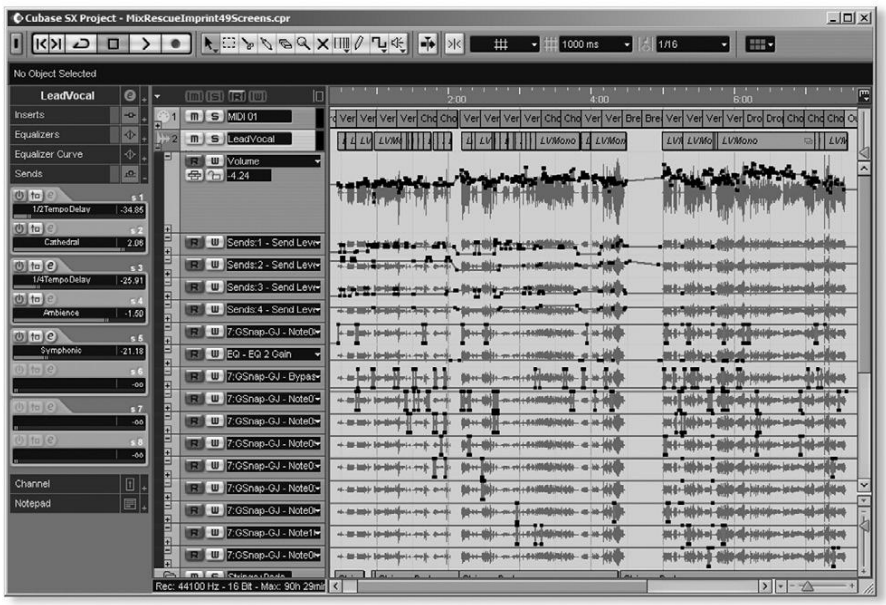

FIGURE 6.7
This screenshot from one of my *Sound on Sound* "Mix Rescue" remixes shows all the vocal automation, including several tracks dedicated to making the automatic pitch corrector (GVST's GSnap) function more musically.

is to make the effort to match the speed and strength of the plug-in's pitch shifting to the needs of each audio region. To be specific, you should try to keep the correction as slow and gentle as possible, while still shifting the note pitch centers to where they need to be. Bear in mind that automatic pitch correctors tend to sound best when they don't have to shift anything too far, so you may be able to improve the audio quality by manually dealing with larger shifts offline in the first instance.

Another group of software applications tackle the problem of pitch detection by first scanning the audio in advance in order to analyze the audio data more intelligently offline. This not only leads to fewer pitch-tracking errors (and the facility to correct any that do still occur), but it also allows for more sophisticated user-guided, note-by-note processing. The latest versions of Antares Auto-Tune, Celemony Melodyne Editor, and Waves Tune all provide this kind of functionality, and this is the kind of processing that I would recommend to those who want the fastest route to musical processing. In particular, their ability to adjust note pitch-center offsets independently of shorter-term pitch fluctuations is brilliant for delivering sensitive correction only where it's actually beneficial.

Pitch-correction plug-ins have all sorts of meters and graphs telling you what they think they're doing, but it's normal to disagree with them from time to time if the underlying pitch-tracking algorithm momentarily slips on a banana skin.

Irrespective of the kind of pitch-correction tools you use, the crucial thing to remember is not to trust your eyes more than your ears—just as when editing timing, in fact. Pitch-correction plug-ins have all sorts of meters and graphs telling you what they think they're doing, but it's normal to disagree with them from time to time—for instance, if the underlying pitch-tracking algorithm momentarily slips on a banana skin. "A lot of people do the tuning offline, so to speak, as a standalone operation, as opposed to doing it while listening to the track," complains Joe Barresi. "[I] make the pitch adjustments while the rest of the track is playing. What I'm doing is listening for the pitch cues that the track can give you. It sounds worse if the vocal is in perfect pitch but the track isn't."[7]

Regular breaks during the correction process are very important, suggests Marius de Vries: "It's definitely possible to use the technology to stamp the life out of musical performances…. You sit back and listen to it and it's totally flat. In the beginning you often won't admit this to yourself, because you cannot believe that you just wasted a day making something worse, but the fact is there…. The only remedy is to keep walking out of the room, and keep walking back in again once you have a clear head."[8] I'd also add that it's worth playing the whole song again from the start once you've finished all your corrections on a given part, listening back with the plug-in window closed, so that you're absolutely sure that things sound right in context over the longer term—whatever any graphical displays might be trying to tell you.

6.4 THE NEED FOR SPEED

Once you get the hang of the general points I've covered in this chapter, decent-quality tuning and timing correction isn't actually that difficult to achieve, even where budgetary restrictions preclude the use of additional third-party software. However, the reality of small-studio life is that there can be quite a lot of day-to-day audio patchup work. So my last piece of advice is to do everything you can to speed up the process, whether that means investing in dedicated additional software or simply configuring your DAW system's keyboard shortcuts for minimum mouse exercise. As vital as it is to commercial record production, corrective audio editing is never going to be a white-knuckle ride, so at least do your best to make it mercifully brief.

CUT TO THE CHASE

- Most small-studio productions suffer from insufficient attention to detail where tuning and timing are concerned. Correcting them is tedious, but few budget productions will compete commercially otherwise.
- To avoid killing the groove, locate a single instrument (often the drums) as a timing reference, refine its timing if necessary, and then edit sloppily performed parts to agree with that. Rhythm tracks usually benefit from the tightest timing, especially double-tracked parts destined to be panned across the stereo image. Lead vocals can also impact on the groove, so be sure to edit them too if necessary. Don't forget to pay attention to where notes end—especially with bass instruments. To achieve best results with timing edits, listen to at least two bars preceding the section under scrutiny.
- Simple cut-and-crossfade audio editing is quick and effective for most timing-correction tasks, as long as you learn where to hide your edit points: in silences, in noisy signals, or at premasked and masked locations. Where concealment is impossible, try a matched-waveform edit. Time stretching has useful gap-filling applications, but the bundled processing within most DAW systems is more than adequate for this, so there's little justification for additional software investment.
- Tuning problems are as much a mix problem as a musical problem. The quality of widely available pitch-shifting algorithms is such that there is no need to have unnatural processing side effects on display if you avoid processing noisy, transient-rich, and polyphonic sounds where possible. Focus on lining up the perceived pitch centers of notes, while treating shorter-term musical pitch fluctuations with kid gloves; and adjust the pitch-processing algorithm to suit the audio being processed. If you're using specialized automatic pitch-correction software, provide it with as much human input as you can to retain musicality. Algorithms that work by prescanning the audio usually offer more opportunity for user direction, so they typically give more musical results.
- Although using visual guides can speed up your work, the cardinal rule with any corrective edits is that you should never trust your eyes over your ears.

Take frequent breaks while working, and audition the whole part top to tail once you've finished working to gain perspective on the success of your processing decisions. With timing edits, it may prove desirable to apply overall timing offsets to whole tracks at this stage, although these offsets may need refinement later in the mix as the final balance comes together.

Assignment

- Correct any timing or tuning anomalies in your mix project, but take care not to compromise the musicality of the original recordings.
- If you do a lot of pitch-correction work, then invest in specialized software to improve your work rate and sound quality.
- Editing is unavoidably tedious work, so find keyboard commands to speed it up wherever possible.

www.cambridge-mt.com/ms-ch6.htm

CHAPTER 7

Comping and Arrangement

Although this isn't a book about recording, there are two specific aspects of the tracking process that I want to discuss briefly in this chapter because they so regularly shoot small-studio mixes in the foot: comping and musical arrangement. Both these factors are important for almost every type of production, regardless of the budget level, yet small-scale projects almost universally undervalue their significance. If your comping and arrangement are good, then mixing is a lot less like hard work; but if they're not, you can be left with an underwhelming record no matter how top-drawer the musical material or your mixdown skills. "I think there's too much dependence on the final mix," says Tony Platt, "and I don't think enough thought and preparation goes into the processes that go on beforehand.… The mix should begin the moment you start recording, by moving things around, balancing, blending, and creating sounds. It means you're creating a coherent piece of work rather than recording a whole series of events any old way and then fixing it in the mix."[1] So let's try to head off what tend to be the most common shortcomings.

7.1 COMPING

Comping is very much standard practice in professional productions. It involves recording a given musical part more than once so that the best bits of each take can be edited together to build up the best composite performance. Clearly this is more time-consuming than just capturing a one-take wonder, so it's not something that small-studio operators need do for every part, but any important solo/hook parts will usually benefit from judicious comping. If nothing else, lead vocals should almost always be built up via comping, as I've yet to hear a one-take lead vocal on any amateur mix that couldn't have been improved in this way. Steve Lipson notes, "[Sometimes] you need to take a really close look at things, like when you're comping vocals. Those are situations where you really can't let any mistakes go, and it can really pay off."[2]

Lead vocals should almost always be built up via comping, as I've yet to hear a one-take lead vocal on any amateur mix that couldn't have been improved in this way.

Mixing Secrets for the Small Studio.

Of course, a comp is only as good as the raw takes from which it's built, and different engineers have different preferences here which range, roughly speaking, between two extremes:

- Recording numerous full "top to tail" takes, and then editing between them after the recording session. The upside of this method is that the musical flow through the part will tend to be more convincing, and you're more likely to maximize the number of unrepeatable little expressive moments you catch for posterity. The downside, however, is that working this way demands much more concentration and endurance from the musician, so in practice you'll only really get superior results with the most talented and experienced practitioners. Less accomplished performers will simply tire themselves out before they nail the more challenging sections and won't be able to maintain enough tonal consistency to allow seamless editing between takes. And if you have to do loads of full-length takes to get all the performance magic you need, then you'll be faced with the mother of all editing jobs to comp them into a final master take. "It takes a long time," said Pascal Gabriel when talking about his work on Dido's "Here with Me." "Out of a two-day writing session we probably spent about three to four hours just comping."[3]
- Recording a single track one phrase at a time, punching in on it repeatedly to create a "patchwork" composite performance. The advantages here are that you concentrate time and effort more efficiently on the sections that need it most; the performer only has to concentrate on perfecting one small section at a time, so it's easier for less experienced musicians; the tone should remain fairly consistent between takes for editing purposes, even as the performer begins to tire; and the comping will effectively be achieved as you go along, so there's no editing to be done later. The disadvantages of this approach, however, are that you lose some of the sense of musical flow between sections of the track, so the final master take can end up feeling slightly disjointed emotionally, and there's also the likelihood that you'll lose nice little performance corners during the real-time patchup process.

Neither of these extremes tend to work particularly well in typical small-studio environments, however, so my recommendation is to chart a course somewhere between them. In most instances, deal with the production one musical section at a time (individual verses, choruses, and so forth) and record a fixed number of section-length takes for each section. If you want to improve any individual take with the odd punch-in, go ahead, but don't get bogged down replacing anything smaller than a whole phrase. I'd say that three takes per section should be considered a minimum if you're going to bother comping at all, and if you're looking for a professional-grade lead-vocal comp, then you'd be well-advised to increase that number to more like eight. Record the takes one after the other until you reach your predetermined limit. By this point your performer will have effectively rehearsed the track several times and

will almost certainly be performing better than they were to start with, so now go back through and listen to each take in turn, and rerecord any which don't feel up to the standard of the rest. Do keep an ear out for magic moments within an otherwise lackluster take, though, because you don't want to erase any of those. What I usually find is that a player or singer will only really "inhabit" the musical part properly after a few takes of each section, so I usually end up rerecording at least the first few takes. That said, the freshness of the first take does sometimes harbor a couple of fantastic little inflections, so listen to that particularly carefully before ditching it. "The bits that most people remember," says Pascal Gabriel, "are the little quirky bits you'll find from a take where [Dido] didn't really feel she was confident with the song, but she let rip a bit by accident, and those bits are sometimes worth their weight in gold."[4]

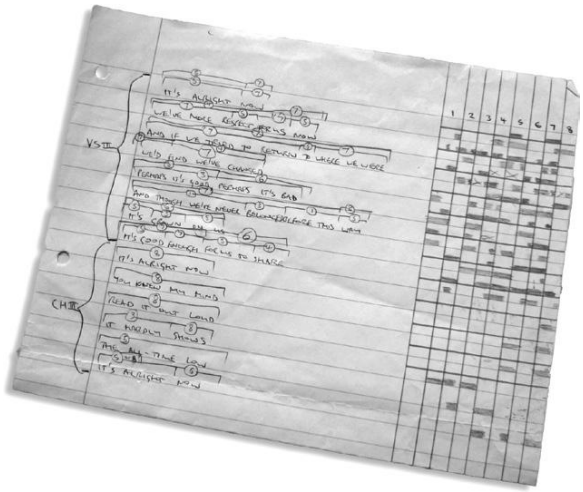

FIGURE 7.1
An example of a real comp sheet for a lead vocal during a song's verse and chorus, and also for a harmony part during the verse. The columns at the right of the page indicate promising sections of the eight vocal takes, whereas final decisions are marked over the lyrics with numbers and brackets.

Once the vocal session has finished, you should be left with a limited number of fairly strong takes from which to choose highlights at your leisure, but at the same time you won't have hours and hours of recordings to toil through. Editing should be straightforward too, because the tone between takes should be quite closely matched. Furthermore, you won't have put performers through the wringer too much, because they've only had to concentrate on a small section at any one time (perhaps over several separate sessions), but the section-length takes will usually lend the comped performance a reasonable sense of musical fluidity.

When it comes to selecting the best bits of your assembled takes, the time-honored route is to draw up a "comp sheet" on lined paper. Each horizontal line of the sheet corresponds to a musical phrase—in the case of an instrument part, you might indicate which is which with bar numbers, whereas in the case of a lead vocal, you'd write the phrase's lyrics on each line. A series of columns down the right-hand side of the sheet can then be used to write comments relating to the quality of each take for any given phrase. What I normally do is listen to each take all the way through marking any absolutely "must-have" sections and then try to build as much of the part as possible from those snippets. Only when a phrase doesn't have a standout rendition on one of the takes will I then get into comparing all the individual takes of a given phrase—it might seem more efficient to do phrase-by-phrase comparisons like this straight away, but I find that this tends to encourage more hopping between the takes, so the composite performance doesn't make as much musical sense overall.

VOCAL COMPING: WHAT TO LISTEN FOR

Choosing vocal takes is a subjective pastime, and it can be tricky for those just starting out to get a handle on what to listen for. Mike Clink has some good general-purpose advice: "The number one pointer that I can give is to visualize what you're trying to say. I'm trying to see in my mind if the singer is telling me the same story as what he's written down lyrically.... A vocal can't just lay there in a linear manner; it's got to take you up and down, depending on what the lyrics are trying to say... If you can look at your lyrics and listen to your vocals and say 'it's telling me a story, and it's taking me on a journey,' then you've accomplished a great performance."[5]

"I'd always look for a vocal performance that had feel over pitch perfection," adds Serban Ghenea. "[Even if it's difficult to correct, it's] worth it because you get to keep the takes that have the best feel, and the best vibe."[6] A final perceptive tip for working with rappers comes courtesy of Mark Ronson: "There's the rhythm and the clarity, but there's also a melody in rap. With the really good rappers, even if it's not a song, the melody takes you on a bit of a trip, and I always try to find the little bits in the vocal that have that melody. Sometimes that's almost sing-songy, but not."[7]

Once the comp sheet is completed, the nuts and bolts of the editing shouldn't present many fresh problems if you've already got a grasp of the editing techniques discussed in Chapter 6. However, you should take every precaution to avoid the edits becoming audible or else you'll risk defeating the whole point of doing the comp in the first place—few listeners are likely to find a performance compelling if it sounds like it's been stitched together by Dr. Frankenstein's intern. For this reason, a number of edits may need considerable fine-tuning before they pass muster, and you may have to adjust the level, pitch, or timing of the odd audio slice to get all the pieces of audio to fit together naturally.

7.2 BREATHING LIFE INTO YOUR ARRANGEMENT

If you think of the basic musical material in a production as consisting of lyrics, melodies, and chords, then the arrangement of that musical material constitutes pretty much all the other aspects of the track—things like the specific instrumentation at different points, the exact distribution of the harmony notes across the pitch range, the particular performance figurations used by each different performer, and the precise rhythmic configuration of the music's groove. If the arrangement works, then it'll play into the hands of the mix engineer, whereas if it fights the music, then there's not much that can counteract that musical ineffectiveness at mixdown. "When you understand arrangement, mixing is easy," says Humberto Gatica. "You know what you're looking for. Otherwise you're fishing."[8] Andy Johns is one of the many other producers who share his opinion: "The way I really learned about music is through mixing, because if the bass part is wrong, how can you hold up the bottom end? So you learn how to make the bass player play the right part so you can actually mix."[9]

Clear Out the Clutter

The first common problem that I hear in small-studio productions is that there are simply too many parts going on all the time. Not only does such a situation make it difficult to make all the parts clearly audible in the mix, but it also presents a musical problem—although a full production may appear satisfying at the outset, a lack of light and shade in the arrangement will quickly make it seem bland, and the listener will stop paying attention.

The first thing that can help is to bear in mind this simple piece of advice from Jack Joseph Puig (echoing an almost identical comment from Wyclef Jean[10]): "You have to consider the fact that the ear can process only three things at once. When you get to the fourth thing, the attention drops away somewhere."[11] So listen to your track and try to isolate which three elements should be the focus of attention at every point in the music. Any part that isn't one of those main components is a candidate for pruning, so see if you can chop out some sections of it. Let's say that there's an acoustic guitar part strumming away throughout your track; perhaps you should consider taking it out of your verses so that more important drum, piano, and vocal parts can come through more clearly. Save it for adding some extra textural interest during the choruses.

Although a full production may appear satisfying at the outset, a lack of light and shade in the arrangement will quickly make it seem bland, and the listener will stop paying attention.

CLASHING TIMBRES

A common problem in less accomplished arrangements is when two sounds clash in some way, usually by operating in the same register. "The aim is to get all the sounds working together so you don't get any nasty surprises at the mixing stage," explains Steve Power, sharpening his scythe. "If two things sound crap together, you probably shouldn't be trying to [fix that] at the mix. If they sound crap, just don't put them together in the first place, because you probably won't rescue them."[12] Eliminating a part completely may not be an option, however, in which case this Fleetwood Mac anecdote may provide some inspiration. "When we were recording *Rumours*," recalls Ken Caillat, "[keyboardist Christine McVie] would ask, 'How does everything sound, Ken?'… Sometimes I'd say 'You know Chris, I'm having trouble hearing the keyboard and the guitar.' The first time I said that I didn't really know what I meant, but she said 'Oh, yeah, you're right Ken, we're playing in the same register. Why don't I invert the keyboard down a third and get out of [the guitarist]'s way?' Which is what she did, and it worked brilliantly!"[13]

Or perhaps you might edit that guitar part so that it only plays at those moments when one of the other parts doesn't really demand the listener's full attention—perhaps while the singer's taking a breath or the pianist is indulging

in some uninspired vamping. This effectively substitutes the guitar for one of your three main parts just for a moment, without robbing attention from them during the bits where they're more important. Most parts in musical arrangements are only playing something interesting some of the time, so if you try to cut out as much of the less interesting stuff as possible, then you will generate some gaps in the texture for other subordinate parts to have their place in the sun. Just this simple arrangement trick can completely transform the sound of a production, because it focuses the listener's ear on what's interesting in the music. There will of course be a limit to how ruthlessly you can perforate the arrangement before the musical flow begins to suffer, but most small-studio producers don't get anywhere near this point before they throw in the towel. As Hugh Padgham observes, "It often takes a lot of effort to have less rather than more. I actually spend more time pruning stuff down than adding things. Doing so can often require a musician to learn or evolve an altogether different part to be played, so what was two tracks is now one track. Every song is different, but I'm always looking for ways to simplify and reduce."[14]

Another thing to consider is whether you might be able to differentiate the sections of your arrangement more effectively by your choice of parts. For example, the first four sections of a common song structure might be intro, verse 1, bridge 1, and chorus 1 (as shown in Figure 7.2), but it's going to be difficult to create any musical development if all of those sections are backed by the same drums, bass guitar, strummed acoustic guitar, and rippling piano figuration. Why not start with acoustic guitar, bass guitar, and drums, but then drop out the acoustic when the lead vocal enters in the chorus? That'll make the vocal appear more interesting and give you the opportunity to differentiate the bridge section by adding the guitar back in again. Then you can leave the piano entry to provide a lift into the chorus. It's ridiculously easy to experiment with your DAW's mute buttons to see what might work here, so there's no excuse for letting the grass grow under your feet. In a similar vein, let's assume that verse 2, bridge 2, and chorus 2 are next in line. It's daft to just repeat the same arrangement, when you can intensify what went before. So maybe let the piano continue from the chorus 1 into verse 2, but drop out the acoustic guitar. The guitar could then join the piano as before to create a thicker bridge 2 texture, and then you might double-track the acoustic guitar or add other parts (such as backing vocals or tambourine) to give another step up into chorus 2.

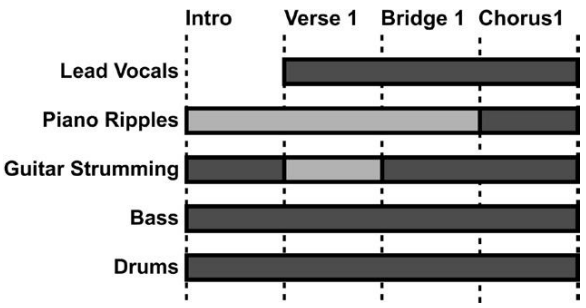

FIGURE 7.2
A static arrangement can undermine the structure of your song, so try to prune your parts to emphasize each section with some kind of arrangement change. An example of this would be to mute the lighter-shaded parts in the hypothetical arrangement shown here.

To put all that in general terms, in a lot of cases in commercial music you want to have enough repetition in the arrangement that the music is easily comprehensible to the general public. But you also want to continually demand renewed attention by varying the arrangement slightly in each section, as well as giving some sense of an emotional buildup through the entire production. That

may sound like a bit of a tall order, but it can often involve no more effort than poking the odd mute button.

Marius de Vries offers another interesting angle: adjusting the musical timeline. "What's often most boring about contemporary songwriting," he points out, "is

THE DROP CHORUS

It's not uncommon to have a double chorus at the end of a typical commercial song's structure in order to really hammer home the main hooks. This is a sound idea in principle, but it does present one arrangement difficulty: if you give the first of these final choruses a really full arrangement to let everyone know it's arrived, then it becomes difficult to sustain that excitement through into the second of the final choruses, because there's no room left in the mix for additional parts. A lot of productions address this problem by using more impassioned vocal parts and adlibs, but I've always felt that there's only so far you can go with that approach, and that there's more mileage to be had in using more imaginative arrangement alternatives instead.

FIGURE 7.3
Some albums contain excellent examples of drop choruses, such as Jordin Sparks's *Battlefield*, Dido's *No Angel*, and Anastacia's *Not That Kind*.

One trick that I just can't resist mentioning is what is often known as a "drop chorus." You build up to the first of the final choruses as if it's going to be huge, but then wrong-foot the listener by muting loads of the important parts when it does arrive, typically leaving just the lead vocals and some kind of stripped-back rhythm section. If you judge it just right, the listener effectively gets two chorus pay-offs at the end of the song, instead of only one. The first of the final choruses uses the element of surprise to make its impact (as well as focusing attention clearly on the lead vocal, which is rarely a bad thing!), whereas the subsequent chorus delivers that all-guns-blazing fullness that everyone was expecting first time around. If you've not consciously noticed any song using a drop chorus before, I encourage you to keep your ears open for it, because it can be extremely effective—Jordin Sparks's "Battlefield," Dido's "Hunter," and Anastacia's "I'm Outta Love" are just some of the hit songs that have made dramatic use of the concept. What's particularly great about a drop chorus from a mix engineer's perspective is that you can often create one entirely at the mixdown stage if you want, without doing any further tracking at all—all you'll usually need is a bit of muting and rebalancing.

the way that everything is in bars of four and blocks of eight, which is arguably driven by the way the software is designed. Good creative songwriters will always be searching for ways to jettison a beat in a bar or add another couple of beats before releasing the chorus, for example."[15] Even if you're mixing other people's material, it's surprising how much you can still experiment with these kinds of changes using simple editing and time/pitch-manipulation tools.

Adding Detail

The other main shortcoming of ineffective small-studio arrangements is that they don't have enough useful detail in them to sustain interest. Again, there's usually too much repetition occurring. A useful little rule of thumb that can help here is this: if you want something to retain the listener's attention, then avoid playing the same thing more than three times in a row. By the time a given riff reaches its fourth iteration, the listener will start to drift, so if you want to keep people focused, then you need to introduce some variation, usually in the form of some kind of fill. So if your drum part is a one-bar loop and your song works in eight-bar phrases, then it makes sense to perhaps edit the loop a bit during bars 4 and 8 to create fills, providing some variation to maintain awareness of the loop among your audience. How exactly you do the fill is up to you, but remember (as I've already mentioned) that the best location for a fill on any given part is usually where other parts are less interesting. If you feel you want to layer more than one fill simultaneously, then consider putting them in different pitch registers—so combine a bass fill with a snare fill, for example.

If that little "three in a row" test reveals some of your parts to be wanting and a bit of audio editing can't generate anything significantly less repetitive, then it's a signal that you probably want to get other musicians involved—a good performer will intuitively create fills and variations at appropriate points in the arrangement. Even if you're stuck on a Hebridean island with nothing but seagulls for company, that needn't rain on your parade, because you can order made-to-measure live overdubs by talented session musicians over the Internet these days—services such as eSessions, Studio Pros, and The Missing Track can all record a wide selection of live instruments over your backing tracks for surprisingly little money.

If you're producing seriously chart-oriented music, make sure that there's some interesting little feature happening every three to five seconds to cater for the attention-span of young media-bombarded music fans.

If you're producing seriously chart-oriented music, then another useful guideline is to make sure that there's some interesting little feature happening every three to five seconds—perhaps a particularly cool backing-vocal lick, a nifty drum fill, a synth hook, a little spot-effect on

the lead vocal, some kind of wacky SFX sample—the possibilities are endless. Chart music needs to command attention continuously if it's to cater for the attention span of young, media-bombarded music fans. One other little tip is to consider the bass part of your track not just as a part of the harmony, but more as a second melody line. Brian Malouf says, "Generally, I try to have the bass and vocals be the counterpoint to each other. I look at those two elements as the central melodic components."[16] The more singable you make the bass, the more forward momentum it'll tend to add to the production. The classic example I always think of here is Abba's "Money Money Money," but if you listen to a lot of good arrangements, you'll frequently hear little melodic fragments cropping up all over their bass parts.

CUT TO THE CHASE

- Whatever style of music you're working in, if any of your most important parts are recordings of live performers, you should consider recording several takes and then comping together the best bits. Lead vocal parts in particular are routinely comped from many different takes in a commercial context. How you go about recording and compiling the takes can have an enormous impact on the quality of the outcome, so don't be afraid to adapt your techniques to suit the particular part and performer you're working with.
- Reducing clutter in an arrangement not only makes it easier to mix, but also helps the production appear more varied to the listener. One general-purpose approach is to think in terms of restricting yourself to a maximum of three main points of interest at any one time and then to weed out as many parts as you can that don't form part of this main trio—but do bear in mind that the relative appeal of individual instruments can vary from moment to moment, so the identity of the three primary parts may change frequently. Also try to alter the arrangement to give each new musical section its own sound—this will better maintain the listener's interest and support the buildup of the production as a whole.
- Boring arrangements usually suffer from too much repetition, so consider adding some kind of fill if any part plays the same thing more than three times in a row. If you can aim a fill in one part to grab attention from a less interesting moment in another part, then that's a bonus! If you're working in chart-oriented styles, then try to provide some new musical or arrangement diversion every three to five seconds to keep the listener riveted to the radio. Treating the bass line as a second melody is also surprisingly effective in improving musical momentum.
- If your song ends with a double chorus, but the second of the choruses seems like it's treading water, experiment with your mute buttons to see if a drop chorus might bail you out.

Assignment

- If you haven't already comped the most important live recorded part in your chosen mix, then do it now. Compare the comped version to your original one-take wonder to get a sense of why comping is such standard practice on professional productions, especially for lead vocals. Bearing in mind the time-demands of the comping process, decide if any further parts might warrant similar treatment.
- Go through your arrangement and try to identify which are the three most interesting things happening at any given time; then see how far you can thin out the remaining parts behind them to declutter the sound. Check for sections of your track that share similar arrangements. If you find any such sections, ask yourself if you might be able to support the music better by altering one of them, either to provide variety for the listener or to enhance the overall buildup of the production.
- Hunt down any musical parts that play more than three times in a row, and see if a well-placed fill might make them more engaging. Look out for any opportunities to give the bass line some melodic function, and decide whether that helps drive the music along better.

www.cambridge-mt.com/ms-ch7.htm

PART 3
Balance

Your mix project is neatly laid out. Timing and tuning gremlins have been banished. The arrangement has been honed. All your faders are pulled down, and you're ready to mix.

Now what?

That unassuming little question is essentially the crux of what mixing technique is all about, and answering it with conviction at every stage of the process is what foxes a lot of small-studio operators struggling to get commercial-level results. The reason so many low-budget producers can't confidently answer the persistent "now what?" is that they don't have an overarching game plan to provide solid rationale for their processing decisions.

Observing seasoned mix engineers at work is very deceptive in this regard, because they'll often seemingly skip at random between different mixing tasks. In reality they have developed such an intuitive grasp of their own mixing workflow that they are able to respond freely to mix issues as they arise without the risk of losing the underlying plot. As such, I don't think it's very

Observing seasoned mix engineers at work is very deceptive, because they'll often seemingly skip at random between different tasks. In reality they have developed such an intuitive grasp of mixing workflow that they can respond freely to mix issues as they arise without losing the underlying plot.

productive for anyone to learn to mix in this way at the outset, because it's not the apparent randomness that's important; it's the underlying plot!

Therefore, Part 3 of this book unravels the nonlinear real-world mixing procedure into more of a bite-sized, step-by-step process so that's it's easier to digest and understand. This inevitably means that I'll be structuring the mix workflow more rigidly than an experienced practitioner would, but by doing this I hope to clarify the thought processes that lie behind even the most intimidating and complex mix techniques. Only once you're thinking logically about the mix process can you be sure you're making decisions for solid reasons—at which point you'll be ready to progress to more a freestyle professional approach with confidence.

Consider the analogy of learning to drive. You don't go straight out on the motorway for your first driving lesson, even though motorway driving is a big part of what most drivers normally do. You first have to learn a variety of potted techniques in isolation: operating the indicators, controlling the clutch, manipulating the steering wheel, swearing at couriers, and so on. It's only once you have the experience to juggle all these tasks subconsciously that you can cruise the fast lane at 80mph while listening to the football commentary and furtively dunking a chicken nugget.

The first task when mixing is finding a balance—at once the simplest concept to express and the most challenging to implement. If you can get every important sound clearly audible and at an appropriate level throughout the duration of the mix, then the battle's mostly won. As John Leckie has noted, "Very often, rough mixes are used on an album. The idea of a rough mix is simply to hear what you've got, so you make your balance so that you can hear everything. And, really, that's all you want out of a finished record—to hear everything. The worst thing is things getting obscured. People will turn the music off when they can't hear everything."[1]

CHAPTER 8

Building the Raw Balance

8.1 ORDER, ORDER!

The order in which you mix the different instruments and sections within your production can have huge consequences for the final sonics, so the first way you can make sense of the mixing process is by working out where to start.

Start with the Most Important Section

Almost every piece of music worth the name has sections that can be considered emotional high points, and a series of drop-downs and buildups connecting them. Your job at the mix is to support this natural ebb and flow (sometimes called the "long-term mix dynamics"), while reserving the greatest musical intensity for the production's most climactic section. This is something that amateur engineers have great difficulty with, the most common complaint being that the verses sound bigger than the choruses. Fortunately, the solution to this problem is disarmingly simple: work on your song sections in order of importance.

For instance, let's say you're mixing a song with this simple pop structure: verse 1, chorus 1, verse 2, chorus 2, midsection, chorus 3, chorus 4. If you want chorus 4 to be the biggest-sounding section (not an enormously controversial choice!), then mix that first. Throw everything including the kitchen sink at it so that it sounds as breathtaking as you can possible manage. When you've reached the limits of your mixing capabilities, wipe the sweat from your brow and then see how you can adapt the same mix settings to chorus 3 so that it's just a little less powerful—it should still sound great, but it should also leave room for chorus 4 to sound better. Go back to chorus 2 and chorus 1 with a similar aim. You still want them to sound like choruses, but at the same time you want their impact to be slightly less than any of the following chorus sections.

Once you've defined the subjective importance of each of these choruses, you can set about joining up the dots. Let's say that the midsection of your song is designed to pretty much sustain the energy level of chorus 2; mix

Mixing Secrets for the Small Studio.

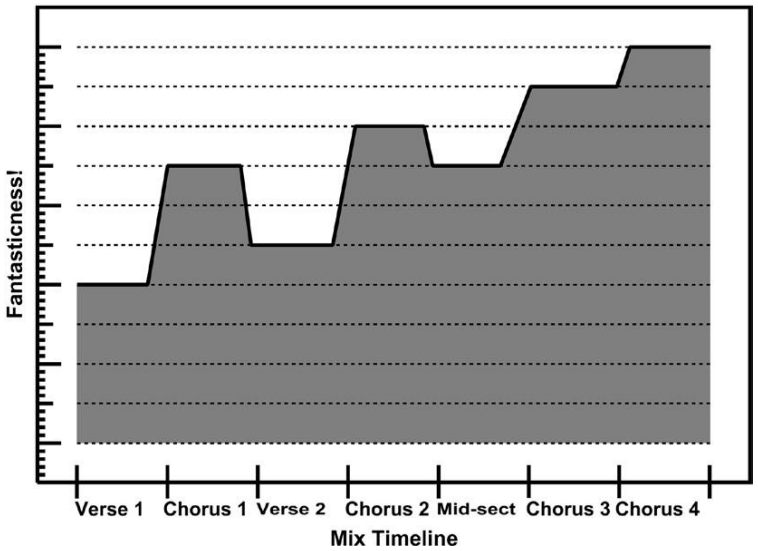

FIGURE 8.1
One possible example of sensible mix dynamics for a typical short song structure. The choruses are more intense than the verses, and each chorus is more intense than the previous one.

that next, perhaps with fairly similar settings as used for the choruses, and then listen how the subjective energy levels change across the section boundaries. If chorus 3 feels like a letdown, then you've pushed the midsection too far. Rein it in a bit. Remember, you can't increase the intensity of chorus 3, otherwise chorus 4 won't sound like a proper climax; and chorus 4 can't get any bigger because you've already maxed out your mixing mojo there. Besides, it's a lot easier to make things sound smaller than bigger anyway. Clearly you'll still want the midsection to sound as good as possible within these limits, but not so good that it undermines the excitement of those final choruses.

Filling in the rest of the blanks in this song structure works just the same: make each section sound as good as you can, but without compromising the musical importance of those sections you've already mixed. Sometimes this means that you end up mixing the beginning of the song last, which stands to reason if everything else is building from there. However, many famous pop songs actually kick off with a modified chorus of some sort, which may not actually be the lowest point in the overall mix dynamics, so mixing your sections in order of importance doesn't mean mixing sections in reverse. It's your responsibility as a mix engineer to tap into the emotional flow of the music in order to decide on the mixing order that will best serve that goal.

Beyond ensuring that a production's climaxes are actually climactic, the hierarchical mixing approach can also highlight very clearly when it's time to use the most underrated of all mix tools: the mute button. If the arrival of our

hypothetical chorus 3 is sounding a bit limp, then a good way to give it more of a boost might be to drop the midsection further by removing parts. Maybe the stereo shakers could bite the dust, or some of the guitar layers, and that sax solo was too cheesy anyway. In more extreme cases, you could even cut out some of the core rhythm section, such as the bass or drum parts. Another common target for cuts is verse 1 if its arrangement is currently identical to that of verse 2—without differentiating the two verses in some way it can be difficult to generate any sense of musical momentum through the first part of the song. At a professional level, there is now an expectation that mix engineers will have this kind of creative input, not only cutting parts but adding new ones too if that's what it takes to achieve results.

FIGURE 8.2
A seriously underrated mix tool: the mute button.

"Mixers make production decisions all the time," comments Jimmy Douglass, "and I add music to people's records when it's not finished or done correctly. In my work as a mixer, I'm a finisher. That's what I do."[2] Jaycen Joshua confirms this approach: "Once you get into the upper echelons of mix engineers, you are being paid for your taste. You're not hired just to make the rough mix sound better. You are the last line of defense creatively, and you're expected to improve the record. It is your job to change things."[3]

One word of warning, though, when working this way: no matter how much you pare down the introduction of your production to allow room for buildup, it's still important from a commercial perspective that it engages the listener straight away. As Rich Costey explains, "If the intro isn't great, it doesn't matter what happens to the rest of the song. That's something I learned a long time ago: if the song is going to work, the intro has to grab you immediately. The importance of this cannot be overstated."[4]

Start with the Most Important Instrument

There are some engineers (such as Bob Clearmountain,[5] Chris Lord-Alge,[6] Jack Joseph Puig,[7] and Cenzo Townshend[8]) who like to throw up all the faders at the same time and work on everything simultaneously from the outset. Although there's no arguing with their results, my view is that it's only really possible to pull off this approach successfully if you've got the kind of mixing confidence born of decades of experience. Those starting out are better served, I think, by following the lead of other equally high-profile names (Mike Clink,[9] Jack Douglas,[10] Andy Johns,[11] Tom Lord-Alge,[12] Alan Parsons,[13] and Tony Visconti,[14] to name but a handful) who prefer to build the mix up in stages, adding instruments progressively until the complete mix is happening.

Within this context, the idea of working in order of importance can also pay dividends, following the lead of Joe Chiccarelli: "It's a matter of finding out what the center of the song is or what makes it tick. Sometimes you build it around the rhythm, sometimes you build it around the vocal."[15] The thing to emphasize here is that it's sonic importance that you want to be evaluating

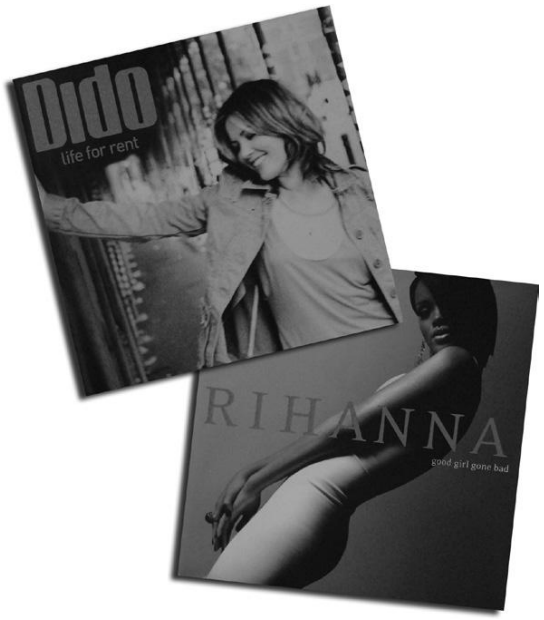

FIGURE 8.3
Two albums that give lead-vocal sonics extremely high priority, often to the detriment of other backing sounds, are Dido's *Life for Rent* and Rihanna's *Good Girl Gone Bad*.

rather than musical importance. A common case in point here is lead vocal parts, which despite their central role in delivering the main melody and lyrics, often aren't as important to the sonic signature of the track as the rhythm section parts. The Sugababes track "Red Dress" is one illustration of this idea—although vocal melody and lyrics are the lifeblood of any pop tune, the most important sonic elements in this track are actually the bass and the drums.

Building up your mix one instrument at a time, in order of importance, means that you avoid the all-too-common problem of running out of space for everything. If a big vocal sound is important to you, then it's a waste of time to spend hours packing your mix full of Godzilla-grade rhythm guitars before you add that in. When you do fade up the vocal, it'll be fighting for its life against the wall of guitars, and you'll have to trash a lot of the guitar processing you've just slaved over. If every part you fade up is less important sonically than the one before, it's much clearer how much space it can safely occupy in the mix, and therefore how full sounding you can afford to make it before it tramples over more important mix elements.

A couple of high-profile commercial productions that provide particularly extreme examples of this principle in action are Dido's "White Flag" and Rihanna's "Umbrella," both of which give over the majority of their mix space to the drums and vocals while savagely carving away some other background parts. In "White Flag" the acoustic guitars are pretty much reduced to a sprinkling of rhythmic pick noise, whereas the synths that arrive in the chorus of "Umbrella" are so hollowed out that they're more an illusion than anything else. Don't get me wrong, I'm not criticizing these mixes; on the contrary, both of them are extremely bold and admirably effective, and what they demonstrate is that mix processing decisions for any given instrument need to take account of the mix space available to it.

There are some other significant advantages of this methodology from a small-studio DAW user's perspective too. Given how easy it now is to put 15 plug-ins on a track in most DAWs, overprocessing is an ever-present hazard for software users. All plug-ins have side effects, and the less processing you can use to solve a problem, the smoother the final mix is generally going to sound. "The thing I've learned," says Justin Niebank, "is to do as little as is humanly possible, because when the computer has to work that much harder—because you've piled on loads of plug-ins, for example—the sound seems to get smaller and smaller."[16] Because the purpose of a lot of mixdown processing is simply to

resolve conflicts between instruments, you'll typically find that each new track you add needs more and more processing to squeeze it into ever smaller pockets in the balance, so later tracks will inevitably suffer more detrimental side effects. By mixing your most important instruments while there's still lots of room to maneuver, they'll end up with less processing, fewer unwanted artifacts, and a more appealing sound.

This way of working can also play into your hands if you're short of processor power, because high-quality plug-ins are usually more CPU-hungry, so you're more likely to use them early on in the mix process when computing power is still fairly thick on the ground. Your most important instruments will therefore be more likely to receive your best-quality plug-ins; and if CPU munch starts restricting your processing options later on, any loss of quality will damage the production sonics less by dint of being relegated to more peripheral parts.

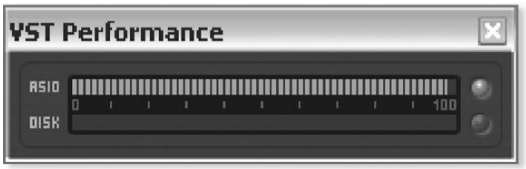

FIGURE 8.4
If you leave mixing important tracks until the last minute, your CPU may max out before you can give them the quality processing they deserve.

If you're just starting out mixing, then it can actually help to write down a list of all your tracks and then number them in order of importance, as a way of disciplining yourself to stick to the plan. It's not always easy to get off the fence and commit yourself, though, especially if you're mixing someone else's work. The relative importance of different instruments in a mix is heavily genre dependent, and you'll often find that an engineer's choice of what to mix first reflects this—so rock producer Mike Clink[17] talks about building up from the drums and bass, for instance, while pop-centric producer Humberto Gatica[18] concentrates on the vocals straight away. As such it's common sense to look to relevant commercial releases for inspiration or seek advice from any engineer you know who specializes in the appropriate style.

A client's rough mix can also be useful in this respect, as clients will often instinctively balance it heavily in favor of the sounds they want to hear. However, do be aware that some self-recording vocalists would rather hear a copy of *Ultimate Manilow* on repeat than listen to their own recorded voice, so take buried vocals with a pinch of salt. Finally, if during your reconnaissance of the project during the mix prep process you've happened on anything dazzling lurking surreptitiously among the multitracks, you might feel that it's worth reversing some of the established stylistic conventions to make this into a much more prominent feature of the track—it's a risk, but at mixdown it's surprising how often fortune favors the brave.

Note that it's entirely normal that the relative importance of certain instruments should vary for different sections of the arrangement, and you need to take account of this when building each new section of the balance. In fact, if you make a point of writing down an instrument ranking for each arrangement section, it can serve as a good indicator as to where multing may be necessary—if the relative importance of any instrument changes dramatically between sections, it's odds on that it'll need multing to allow for different

processing strategies in each instance. Lead vocals, acoustic guitars, and pianos, for example, are often of much more sonic importance during sparser verses than during more densely populated choruses.

Time Is Money

One final general concept to bear in mind with mixing is that time is money. In other words, you should try to allocate the majority of your time to those aspects of the production that are most likely to sell your mix: to the client, to the industry, and to the public. Glen Ballard says, "What we sometimes do in the studio is we spend 90 percent of the time on stuff that's worth about 10 percent—the absolute detail of everything that, at the end of the day, probably doesn't matter."[19] In this respect, spending the same time and effort on every track in your mix is usually misguided, because there are only ever a few aspects of any production that carry the main commercial appeal.

Spending the same time and effort on every track in your mix is usually misguided, because there are only ever a few aspects of any production that carry the main commercial appeal.

What exactly the main hooks in a production are varies a great deal, of course, depending on the style, but it's nonetheless important to be as clear as you can in your mind where the money is. For many mainstream chart styles, the vocals clearly deserve a lion's share of the mixing time (and indeed the mix preparation work), as those elements will command the attention of the widest range of listeners and also carry the main copyright assets: melody and lyrics. It's certainly not at all unusual for vocals to occupy well over half the total production time in pop and R&B, so if you're working in those styles be realistic about what you're up against.

However, in electronic dance styles it may be that the kick, bass, and synth/sample hooks warrant the majority of the attention. Whatever the key tracks in your production, however, the main thing for self-producing small-studio owners is that time is hardly ever on their side, so if you can discipline yourself to abandon some finesse where it's not going to be appreciated, then you'll be able to focus more of your time and energy on polishing those musical features that will really enhance the production's market prospects.

8.2 SIMPLE BALANCING TASKS

Once you've decided which section goes first, you can set up your DAW's transport to loop it for you while you start putting together a basic balance. This isn't a question of merely wiggling faders, though, because there are actually a number of tasks that need attending to at the same time if your level settings are going to have any validity. Stereo panning needs sorting out, for a start, because it's otherwise impossible to judge the effectiveness of the mix in mono (as discussed in Section 2.1). Any comb-filtering effects between multimic recordings

must also be tackled before fader levels can be set reliably, and low-frequency rubbish is best removed as early in the mix process as possible as well, for the reasons mentioned in Section 3.5.

The simplest type of instrument to balance is one that has been recorded to its own mono audio file, without any simultaneously recorded instruments audible in the background (in other words, without "spill," "bleed", "leakage," or whatever else you might happen to call it). With this instrument all you need to do is remove unwanted low frequencies, adjust the pan control, and set the fader level to where it sounds best. Let's look at these three steps in more detail.

High-Pass Filtering

Cleaning up the low end of the track is a job for your first bona fide mix processor: the high-pass filter. What this does is progressively cut away the low end below a user-defined "cut-off" or "corner" frequency. The increasing extent to which low frequencies are attenuated as you progress down the spectrum below the cut-off point is usually expressed in decibels per octave, and some more advanced high-pass filter designs give you some control over this. Although you'll find filter slopes of 72dB/octave or more on some filters, these can have problematic audio side effects, so I stick to comparatively moderate filter slopes of 6, 12, and 18dB/octave (sometimes referred to as 1-, 2-, and 3-pole designs, or first-, second-, and third-order filters) for generic mix cleanup work.

Some filters may offer a Resonance, Bandwidth, or Q control, which creates a narrow-bandwidth level boost at the cut-off frequency, but this feature is irrelevant to the task in hand so it should be set for the minimum boost and can then be ignored. In fact, as a rule, high-pass filters with resonance controls are

FIGURE 8.5
High-pass filtering is occasionally provided as a dedicated plug-in (such as Brainworx's Bx_cleansweep, *right*), but it more commonly forms part of a fully featured equalizer (such as Universal Audio's Cambridge Equalizer, *left*).

more associated with sound-design work than mixing, so they may not be the most transparent-sounding choice for this scenario anyway. You'll normally be better served by the kind of high-pass filter that is provided as part of an equalizer plug-in.

Whatever brand of high-pass filter you use, the key concern is setting its cut-off frequency. Roger Nichols describes a good all-purpose approach here, using the example of an acoustic guitar part: "My rule is: if you can't hear it, you don't need it. Turn on the [filter]…. Start raising the frequency slowly until you can hear a change in the sound of the low end of the acoustic guitar…. Now reduce the frequency by about 15 percent."[20]

Although Nichols targets his filtering primarily at midrange instruments in a professional environment, I think it's better small-studio practice to filter practically everything in the mix in this manner. If you can't rely on your monitoring to inform you whether you're losing important low-frequency information even at your filter's lowest setting, then this is where a spectrum analyzer can provide useful additional insight. And there's another thing I'd add to what Nichols has to say: in my experience the power of a drum's attack can sometimes be softened by high-pass filtering well before the instrument's overall tone changes appreciably. For lower-register instruments such as kicks, snares, and tom-toms in particular, I therefore regularly set the filter's frequency control a little lower and keep my ears open for this easily overlooked processing side effect.

Panning Mono Recordings

Panning simple mono instrument recordings isn't really rocket science, and although some people wax lyrical about trying to create some kind of "natural" stereo panorama, the reality of mainstream record production is that some panning is carried out for solid technical reasons, whereas the rest is done purely on grounds of the mix engineer's personal taste and with scant regard for any notion of stereo realism. For example, most drum-kit recordings usually spread the instruments widely across the stereo picture, but you'd only get that picture if you were actually seated on the drum stool. "If you're standing in front of a drum kit," says Flood, "you hear it in mono—you don't hear the toms panning neatly from left to right in your ears."[21] So you'll have to forgive me if I concentrate on the more objectively justifiable reasons for panning here and leave you to make up your own mind about the rest.

Panning toward the extremes deserves careful thought, as it leads to a 3dB level drop in mono, and might even cause an instrument to disappear completely on the wrong end of a pair of shared earbuds.

From a technical standpoint, there are good reasons for putting the most important musical elements of your mix in the middle of the stereo image, because they'll be the ones that survive best in mono and in situations where only one side of a stereo signal is heard. In most cases, this means that your kick,

snare, bass, and any solo lead vocal will normally inhabit the center. This is particularly important for the bass instruments, as transmitting them from the combined surface area of two woofers gets the best low-frequency efficiency out of any given stereo speaker system. (Off-center bass can also cause problems with pressing your music to vinyl as well, if you're concerned with that.)

When any track is panned off-center, a primary consideration is how its balance will change in mono and single-sided stereo. In this respect, panning toward the extremes deserves careful thought, as it leads to a considerable 3dB level drop in mono and might even cause an instrument to disappear completely on the wrong end of a pair of shared earbuds. However, every producer also has to decide how much it matters to them if a panned sound makes the mix feel lopsided, or distracts the listener's attention sideways from the star attractions at the center of the picture. In practice, what most engineers plump for is a compromise, which is sometimes referred to as opposition panning: if an important mono instrument is panned to one side, then another instrument with a similar musical function is panned to the other side. Opposition panning therefore still provides you with stereo width and interest, but without pulling the overall stereo picture off-center or endangering the musical sense in one-sided stereo playback situations. As such, many panning decisions are more about the width and spacing of opposition-panned instrument groups than about the specific stereo locations of individual sounds.

Setting the Levels

After high-pass filtering and panning, you're ready to set levels. For the first track, all you do is choose a sensible level relative to your mixing system's headroom so you don't risk distortion once the whole arrangement's going. While you could use your channel fader for this purpose, there's something to be said for leaving the fader at the 0dB (unity gain) position and applying most of the required gain either with the mixer channel's input gain control (if it has one) or with a simple gain plug-in inserted into the channel. There are two reasons for this: first, almost all faders have their best control resolution centered around the unity gain position, so it's difficult to make fine level changes if the fader cap's right down at the bottom of its travel; second, the fader's upper end stop can be a pain in the neck if you begin the mix too close to it.

Those starting out with mixing usually underestimate the amount of time and attention they should give to basic level balancing, so resist the temptation to skate over this stage of the mixing process just because it seems so self-evident. Justin Niebank notes, "Balance is way more important than EQ or automation, but people seem to give up half an hour too soon on some mixes. You have to think balance first."[22] Try to zero in on what the level of each new instrument means for the mix as a whole, and be aware that a good balance usually involves some kind of tradeoff: on the one hand, you want the track loud enough that you can hear all the things you really love about the individual sound; on the other hand, you need to keep the track low enough in the mix that is doesn't obscure the merits of more important instruments that are already present.

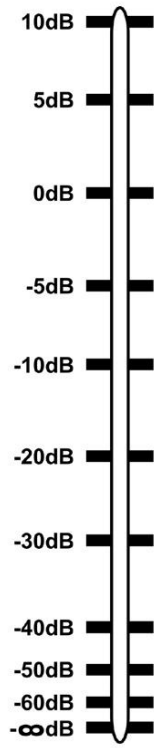

FIGURE 8.6
A typical fader control scale. Fader movements around the unity (0dB) gain position will give smaller and more precise gain changes than fader movements toward the bottom of the level scale.

SOME TRICKS FOR JUDGING FADER LEVELS

If you're having trouble making your mind up about the balance of a particular fader, then there are a few little tricks that can help clarify your thoughts. The first good one comes courtesy of Gus Dudgeon: "I find the quickest way to decide whether something is actually loud enough or not is to run the mix and turn the signal on and off. If you turn it on and you can hear it, but it's not smack bang in your eye, it's probably where it should be.... The best thing to do [if an instrument is too loud] is to take the bloody fader out and start again. Creep it in until you think that's probably where it should be and then try switching it on and off."[23] Mike Stavrou offers another good one: "Instead of focusing all your concentration on the fader being moved, listen to the neighboring instruments instead. While trying to hone the perfect level of the snare drum, for example, do so while listening to the bass drum.... You will quite often find a more decisive point of balance this way. That's not surprising, because when you think about it that's exactly how the listener listens—his (or her) mind is elsewhere until you attract his attention."[24]

"It is always easier to tell when something is wrong than when something is right," says Roger Nichols. "While listening to the mix, one at a time change the levels of each instrument in varying combinations. Turn the bass up 1dB. Does the mix sound worse? If the answer is yes, then turn it back down. Turn the bass down 1dB. Does this sound worse? If the answer is yes, then the original level was right for the bass. Now try [the other instruments].... If you can change the level of an instrument in the mix by a tenth or two-tenths of a decibel and you can hear the change that you made, the mix is getting pretty good."[25] (If that "pretty good" feels like a bit of an understatement, just bear in mind it's coming from a guy with seven Grammies on his mantelpiece!)

A final nifty hint comes from Roey Izhaki. "Take the fader all the way down. Bring it up gradually until the level seems reasonable. Mark the fader position. Take the fader all the way up (or to a point where the instrument is clearly too loud). Bring it down gradually until the level seems reasonable. Mark the fader position. You should now have two marks that set the limits of a level window. Now set the instrument level within this window based on the importance of the instrument."[26] What's particularly useful about this tip is that the size of the "level window" can also provide advanced warning of tracks where further mix processing is required—of which more later in this chapter!

This is where good monitoring practices can really pay off. Use your main and secondary listening systems to gain a multifaceted view of the situation, paying particular attention to your Auratone substitute, and make appropriate use of the low-frequency workarounds discussed in Chapter 3. Keep your ears fresh by switching monitors frequently, varying your listening levels, taking breaks, and seeking other people's opinions. You might even look to your collection of reference productions for guidance, although I'd sound a note of caution there, because it's easy to get bogged down if you try to do detailed referencing early on in the mix process. In practice, I prefer to play through a handful of relevant reference tracks from time to time to calibrate my ears in general, and then

I leave any detailed mix referencing until the end of the mixdown process when the sonics can be compared with commercial productions on a more equal footing.

Although it's a false economy to rush your initial fader settings, you also shouldn't beat yourself up expecting to get a perfect balance straight away, because it takes a lot more than faders to uphold the balance in mainstream commercial genres. Just try to make the best-informed preliminary decisions you can, so that you minimize the amount of processing required further down the line. And if you stay alert during the balancing process, the faders will start to give you valuable tip-offs about which tracks need plug-in assistance.

Listen to Your Faders!

Faders can't talk, but they can communicate nonetheless; it's only once you learn to tune in to their delicate messages that mix processing actually starts to make any real sense. You need to be sensitive to when any fader setting doesn't seem to feel "stable"—in other words, you can't find a static level for that instrument that achieves a balance you're happy with, because wherever you try to set it, it doesn't feel right. If a fader feels unstable like this, then the track probably needs processing to achieve a good balance. This insight might not seem particularly earth shattering if you already know that a good proportion of your tracks will need processing, but what you can already gain from it is the confidence to leave tracks *unprocessed* if they already balance fine at a static fader level. This is a crucial mixing decision in its own right, and, incidentally, it's one that experienced engineers tend to arrive at much more often than rookies do. "One of the smarter things for a mixer is to know when to leave stuff alone," says Jimmy Douglass.[27] Darryl Swann puts it another way: "Taste your food before you salt it. I know so many guys that feel that they're not earning their money unless they're pushing switches and turning knobs. No, you've got to *listen* to your sound. If it needs work, then plug something in."[28]

Leave tracks unprocessed if they already balance fine at a static fader level. This is a crucial mixing decision in its own right, and, incidentally, it's one that experienced engineers tend to arrive at much more often than rookies do.

But just spotting an unstable fader is only the first step, because the precise characteristics of that instability are what provide the essential clues you need to choose the best mix fix. The range of conclusions that it's possible to draw is enormous, and we'll be looking at them in more detail over the next six chapters, but for now here are a few commonplace examples to highlight the kinds of things you should be watching out for during the preliminary balancing stage:

- You fade up your drum loop until the kick seems properly balanced, but you can't hear the snare and hi-hat elements well enough at that level. It feels like you need a separate fader for the kick drum, which is a strong clue that you might want to mult the kick hits to a separate track.

- You try to set the level for a simple eighth-note rhythm-guitar part, but some chords always feel buried in the mix, whereas others pop out unduly. You're constantly feeling the urge to tweak the fader. What your mix is trying to tell you is that the difference in level between the loudest and softest chords is too large—or, to translate that into tech-talk, the dynamic range is too wide. You therefore need processing such as that in Chapters 9 and 10 to control that problem before the fader's going to settle down to a static balance.
- By the time you've faded up your acoustic guitar to the point where you can hear all the nice sparkly string details, the booming of the instrument's main sound hole resonance is making the lowest notes sound bloated. What you want is a separate fader to turn down just that section of the frequency range—which is effectively what the tools discussed in Chapters 11 and 12 provide.

The more sensitive you can make yourself to the way in which a fader feels unstable, the easier it'll be to select the right kind of plug-in for each track without hesitation. For this reason, time spent working on initial fader settings is rarely wasted—the first balance may not singe anyone's eyebrows on its own, but it provides a rock-solid foundation for all the more complex mix-processing decisions later on.

Additional Considerations for Stereo Tracks

All the same things that apply to mono recordings also apply to stereo recordings, as far as initial balancing is concerned, but there are a few additional factors as well. The first thing to address is that some systems record stereo in the form of two mono files (sometimes referred to as "split stereo"), on the understanding that you'll use individual mixer channels to pan the files hard left and right. However, most DAW systems now have perfectly capable stereo mixer channels, which are much more useful for mixing purposes because they allow you to conveniently process both sides of the audio in tandem. If you find yourself presented with a split-stereo recording, make it your business to convert it to a single "interleaved" file before you do anything else. If that's not practical, then buss the left and right mono channels to a communal stereo group channel and do all your balancing and processing there instead.

FIGURE 8.7
A couple of good stereo adjustment and metering utilities, DDMF's StereooeretS (*left*) and Flux's Stereo Tool (*right*).

That said, on occasion you may wish to offset the whole stereo image to one side, narrow the spread, or reverse the sides. In these situations, some control over the panning of the left-channel and right-channel audio streams can

be handy. If you're bussing individual split-stereo channels to a stereo group, then this is no problem. However, things can be trickier if you're working with an interleaved stereo file on a dedicated stereo mixer channel—if that channel sports anything like a pan control, then it probably only adjusts the relative levels of the two audio streams, which is of limited benefit. A few software platforms (Steinberg Cubase, for instance) do actually offer the necessary additional panning controls among the standard channel facilities, but if you can't find any in your own DAW, then there's usually an alternative in the form of a bundled stereo utility plug-in. Still no joy? Then there are plenty of suitable third-party plug-ins around, my current cross-platform favorites being DDMF's pocket-money StereooeretS and Flux's freeware Stereo Tool.

The main concern with stereo files, though, is if phase or polarity issues are causing the sound to change dramatically when you listen back in mono. The sound of any stereo file will always lose something in the transition to mono, so it's unrealistic to expect no change at all, but you do have to decide whether any changes are going to seriously compromise the effectiveness of the balance. If you conclude that the damage calls for remedial measures, then the first thing to check is that one of the two audio streams hasn't inadvertently been polarity-flipped during recording. It's easy to spot this by eye most of the time by comparing the left and right waveforms, particularly for comparatively close-miked stereo files. If your own DAW system doesn't provide the facilities to correct the problem, both the DDMF and Flux utility plug-ins have polarity-inversion buttons that can take care of it.

Even once you've corrected any obvious polarity reversals, phase mismatches between signals arriving at the two microphones may also need addressing if comb filtering blights the mono listening experience. Once again, if you can see a clear delay between the waveforms of the two audio streams, then you may be able to minimize the sonic changes in mono by lining things up more closely, although you should be aware that this procedure may undesirably tamper with the spread of the stereo picture as a consequence. As with timing and pitch edits, do ensure that you evaluate the results of this technique by ear rather than by eye, because the tidiest visual match may not give the most mono-compatible outcome. If you're struggling to find the best phase match between the two audio streams, try inverting the polarity of one of them for a moment, shift the streams against each other until the sound disappears most in mono, and then flip the inverted stream back to normal polarity.

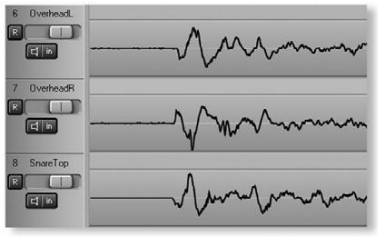

FIGURE 8.8
If your overhead mics have phase and polarity mismatches like this, you'll have real problems with the mono-compatibility of your mix, in addition to a slightly odd stereo picture.

If despite these measures the mono-compatibility problems of your stereo file remain intractable, here's a last-ditch technique that you can still fall back on: pan the two audio streams more centrally, partially summing them to mono in the stereo mix. Obviously this will narrow the stereo image, but it will also partially comb-filter the stereo playback in the same way as the mono rendition, bringing the two tonalities closer together. This may mean that you have some additional processing to do later on to resuscitate the comb-filtered tonality for stereo listeners, but at least the final mix will

translate more reliably to a wide audience. If all else fails, you might also consider ditching one side of the stereo file completely, especially if you already have plenty of stereo interest from other tracks in the production.

So to recap quickly, balancing simple mono or stereo files involves these steps:

- Route both sides of a stereo file through a single mixer channel for processing purposes.
- Optimize the polarity and phase relationships between the left and right sides of a stereo file.
- Use a high-pass filter to remove unwanted low end within the context of the mix.
- Set the instrument's position/spread in the stereo field.
- Balance the instrument against those already in the mix.

8.3 MORE COMPLEX BALANCING TASKS

Multimiked Instruments

It's not at all uncommon in the studio for a single instrument to be recorded using more than one microphone, the idea being that you can thereby adjust the instrument timbre more flexibly at mixdown. It's especially common for guitar amps to be recorded in this way, for instance. Balancing this kind of recording isn't dissimilar to balancing a stereo file:

- Route the individual mic channels to a single mixer channel for processing purposes.
- Optimize the polarity and phase relationships between the different microphone signals.
- Use a high-pass filter to remove unwanted low end within the context of the mix.
- Set each microphone's position in the stereo field.
- Balance the instrument against those already in the mix.

The chief difference in this case, though, is that the second step isn't necessarily just about technical problem solving, because it also has enormous potential for creative tonal shaping. The simplest option is to match the polarity and phase of all the mics as closely as possible (using polarity inversion and time shifts as before), such that you can combine their characteristics fairly freely without any real difficulties arising from unwanted comb filtering. Many great electric guitar sounds have been made in this way, so I'm not going to knock it, but there's nonetheless a whole lot more you can do here if you choose to exploit phase-cancellation effects rather than trying to minimize them.

The simplest option is to match the polarity and phase of all the mics as closely as possible, but there's a whole lot more you can do if you choose to exploit phase-cancellation effects rather than trying to minimize them.

Let's assume for the moment that we're starting from the polarity/phase-matched setup. If you flip the polarity of one of the mics, then rather than adding its characteristics to the composite sound, it'll subtract them. Fading that mic all the way up might unacceptably hollow out the sound of your instrument, but at lower mix levels there will be all sorts of new tonal variations on offer. Likewise, artificially delaying one of the microphone signals a few milliseconds against the others will cause comb filtering between them. Again, you might not want to have the delayed mic too high up in the balance, but at lower levels the milder frequency-response undulations can instantly yield radical tonal changes that would take all day to achieve with traditional mix processors. The best bit about it all is that sculpting your tone with polarity and phase relationships is free of DSP processing artifacts. You'll get big tonal changes with no unwanted side effects—everyone's a winner! Numerous high-profile producers (Ken Caillat,[29] Neil Dorfsman,[30] and Clarke Schleicher,[31] to name just a handful) have professed a fondness for this tone-mangling technique. "When you build a mix," says Jack Douglas, "the real nightmare is when you put something up and the only way you can hear it is by blasting it…. If you do things like this I guarantee you that as soon as you put the sound in the mix, it will be there…. Not only that, it won't wipe out everything else in the mix, because it will have such a separate and distinct character."[32]

The option to pan the individual mics of a multimiked instrument to different positions in the stereo field can be another advantage of this recording method. Simple mono recordings have a habit of sounding much narrower than the instrument itself does when you're seated in front of it, so being able to slightly spread the image of a multimiked acoustic guitar or piano, for instance, can help things seem more natural and real, and it can fill up more of the stereo field in a sparse arrangement. In band recordings, double-tracked electric guitar parts can be made to sound bigger and more impressive in stereo in the same way, as well as more homogeneous if the individual guitar images are actually overlapped to some extent.

You can balance an instrument recording comprising mic and DI signals in much the same say as a straight multimiked recording, because the polarity/phase relationship between the signals creates similar problems and opportunities. However, the extra advantages of a DI recording are that, in the case of electric guitars, it can be reamped to flesh out or even reinvent the sound as required; and that it's pretty much immune to spill from other instruments recorded at the same time, which can be a lifesaver if high levels of spill on the instrument's microphone have rendered that ineffective for adjusting the balance.

FIGURE 8.9
The polarity button. It may not be much to look at, but it's absolutely essential at mixdown, for both technical and creative reasons.

Multimiked Ensembles: Preliminaries

The most complicated tracks to work with are multimiked recordings of groups of voices or instruments: singing guitarists, drum kits, choirs, string groups, live band recordings, orchestras, and many other ensembles. The problem in such situations is that every instrument will be picked up by every microphone to

some extent, and although it's feasible to polarity/phase-match all these different recorded signals for one of the instruments in the ensemble, it's impossible to do so for every instrument at the same time. Imagine a singing guitarist, for example, recorded with one close mic on the singer and one on the guitar. In both close mics you'll have delayed spill from the off-mic sound source, so if you shift the timing of the guitar slightly so that it lines up with the vocal mic's guitar spill, that'll move the guitar mic's vocal spill further out of alignment with the main vocal signal.

The upshot of this no-win situation is that your goal becomes to turn the unavoidable effects of comb filtering to your advantage. Get it wrong and you can comprehensively trash the sound of the recording, making all the instruments sound hollow, harsh, and washed out. Get it right, though, and you can achieve something truly magical: a production where each instrument is enriched by virtue of the contributions from so many mics, and where the combined spill signals naturally blend the ensemble with very little further processing effort. This is where the initial balance can absolutely make or break a record. Unfortunately, this kind of balancing is by no means an exact science, and every engineer has a slightly different way of tackling it. So let's look at some of the options and discuss the pros and cons in each case.

First of all, it's not unusual to have stereo mic pairs within an ensemble setup, and it's usually worth checking those for polarity and phase issues in isolation before you get into dealing with their interactions with other mics in the setup. So, for example, in a typical multimiked drum kit you'd want to individually check the overhead mic and room mic pairs for stereo spread and mono-compatibility. Likewise, if there are any instruments in the ensemble that have been individually multimiked, then deciding on their initial blend and polarity/phase relationship in isolation is also a good idea. In a typical drum set, the snare might have mics both above and below it, for example, or the kick drum might have an inside mic and an outside mic. In both cases, the polarity/phase relationship between the two mics might be adjusted correctively for the best match or creatively to deliver a specific tone. Spike Stent notes, "I'll check whether the kick and snare tracks all line up, and so on. I'm flipping phase all the time. I'm anal about that, because it is essential for getting a really tight mix which sounds big on the radio, cars, laptop, and so on."[33]

Because of all the spill, deciding on a high-pass filter setting for each individual mic becomes more of a gray area, because it depends how much of the low-frequency spill you consider "unwanted" in each case—something that you may only really be able to decide within the context of the whole mix. However, the problem with most normal filters is that, by their very nature, they also adjust the phase response of the filtered signal, which may adversely affect the way the signal phase-cancels against the other mic recordings. My advice is to restrict your use of high-pass filtering here to just removing frequencies well below the range of the lowest instrument; then leave any further filtering until you've had a chance to listen to how all the mics combine in practice (discussed more in Chapter 11).

PHASE ROTATORS AND ALL-PASS FILTERS

I've already discussed polarity inversion and time shifting as two ways of adjusting the combined tonality of multimiked recordings at mixdown. However, there is a lesser-known option, sometimes called "phase rotation," which can also refine the sound in these instances. This is a process that can alter the phase relationship between the different sine-wave components that make up a sound. As with polarity inversion and time shifting, phase rotation has little effect on the individual processed track in isolation. However, if that track is part of a multimic recording, then the phase rotation will change the nature of the phase cancellation between the mics, and therefore the mixed instrument's timbre. As such, phase rotators are useful tools for tonal refinement, although I'd personally recommend experimenting with them only once you've exhausted the possibilities of polarity inversion and time shifting.

Unfortunately, most DAW systems don't currently have a phase-rotation plug-in built in, but there are now several third-party plug-ins that can step into the breach. Voxengo's PHA-979 is a great cross-platform option that is pretty cost effective, and there's also Betabug's VST freeware Phasebug if you're working on a PC. Another option is to use an "all-pass filter," which is also capable of independently adjusting phase response, although in a slightly different way. You can find one in some phaser and EQ plug-ins (check out the freeware Sandford Phaser and Cockos ReaEQ plug-ins, for instance), as well as in the dedicated Littlelabs' IBP Workstation phase-adjustment plug-in.

FIGURE 8.10
A selection of phase adjusters: Betabugs no-frills Phasebug (*left*), Voxengo's fully featured PHA979 (*right*), and Cockos ReaEQ's all-pass filter option.

Panning decisions may be harder to make too, because you need to decide how you pan each instrument's main mic signal in relation to the stereo location of prominent spill from other mics. In this respect, it makes good sense to create a stereo picture that mimics the physical layout of the recording session, because the strongest spill contributions from each instrument should then remain fairly close to the instrument's main mic position in the stereo picture. Where an ensemble recording includes stereo files, matching the images of

these with the panning of any mono mic signals is also fairly commonplace. However, these panning tactics may conflict with the technical panning considerations I talked about in Section 8.2, so there must often be a tradeoff here. For example, a lot of people record drums in such a way that the snare image in the stereo overhead microphones is off center; but you may nonetheless prefer to pan the snare close mics right down the middle for best mono-compatibility, even if this does mean that the tail of each snare seems to veer off slightly to one side.

Building the Ensemble's Balance and Tone

Once these preliminary jobs have been carried out, mute all the ensemble channels and begin reintroducing each mono mic, stereo mic pair, and multimic combination to the mix in turn. As you do this, use your polarity/phase adjustment tools and fader to search for both the best ensemble tone and the best balance within the context of the rest of the mix. If this sounds like a lot of fiddly work, that's because it is, but it's also a part of the mixing process where miracles can happen, so roll up your sleeves and get your hands dirty!

If any mic or mic pair captures a reasonable balance of the whole ensemble, you can often get good results more quickly by using that as the bulk of your sound. Indeed, that's usually my first recommendation for less experienced engineers grappling with multitracked drums.

Again, the order in which you tackle the individual mics and instruments at this stage should reflect their importance sonically. This partly just makes things easier, because the mics covering less important instruments are usually mixed at a lower level, so their spill causes less comb filtering with more important instruments which are already present—you therefore have more freedom to adjust the polarity and phase of the newly added instrument for the sake of its own timbre before causing unacceptable sonic damage to other more critical parts. However, this order of working also clarifies your tonal decisions. For example, let's say you've adjusted the phase of your newly added instrument mic to give that instrument the best tone you can. If you check its impact on the rest of the ensemble by switching it in and out of the mix and discover that it's deteriorating the sonics of some of the previously added instruments, then you don't need a PhD to work out that some aspect of the newly added instrument's sound may have to be sacrificed for the greater good.

There is one complicating factor, however, that can frequently overrule this guideline. If any mic or mic pair has been set up so that it captures a reasonable balance of the whole ensemble, you can often get good results more quickly by using that as the bulk of your ensemble sound: set a level for that first, and then use the remaining signals in the manner of supplementary "spot mics," just for small tone/balance refinements to the main blend. The advantage of this approach is that the spot mics can usually then be used at lower level, thus

reducing the severity of any phase cancellation between their spill contributions and the signal from the dominant overall ensemble mic(s). Indeed, because this method is a little more foolproof in general, it's usually my first recommendation for less experienced small-studio engineers who are grappling with their first experience of a multimiked ensemble—almost always in the form of live multitracked drums. But this approach is by no means just for beginners. Allen Sides, for example, states a definite preference for this method,[34] and Steve Hodge takes a similar line: "I'll start with the overheads first and then add the close-in microphones. That gives the sound some context."[35]

One final issue deserves some discussion for the benefit of those with maverick tendencies. Polarity and phase adjustments provide tremendous freedom for tonal reinvention of multimiked recordings, and in a situation where all the mics are picking up the same instrument, you can pretty much run with that until your legs fall off—no matter what timbre you end up with, you'll still be able to choose whatever level you like for it in the mix. However, it's riskier to take tonal flights of fancy by deliberately emphasizing phase cancellation between different instruments in a multimiked ensemble, because it can interfere with your ability to control the most vital thing of all: their balance. Imagine the example of multimiked drums again: if you deliberately use the rack-tom close mic to cancel unwanted frequencies in that drum's tone, it'll be of limited use to you if you need to fade up one of the tom fills in the final mix.

There's such a lot to think about when balancing multimiked ensembles that it can be tough keeping an overview in your mind, so allow me to quickly recap the procedure:

- Route stereo files and individual multimiked instruments through single mixer channels for processing purposes, and optimize the internal polarity and phase relationships of each group of files in isolation.
- Use high-pass filtering conservatively, removing only unwanted low frequencies comfortably below the range of the lowest instrument.
- Set the position/spread of each mono mic, stereo mic pair, and multimic combination in the stereo field.
- Introduce each mono mic, stereo mic pair, and multimic combination into the mix in turn, adjusting the polarity and phase of each to achieve the best tone and balance for the ensemble as a whole. It's generally sensible to work through the tracks in order of sonic importance, but if the entire ensemble has been captured passably by a single mic or stereo pair, then adding that first as the main bulk of the ensemble sound can often make things easier.

8.4 CASE STUDY: MULTITRACK DRUMS

To clarify how all these balancing guidelines work in practice, let's walk through the balancing process for an imaginary multitrack drum recording as a kind of mental case study. For the sake of argument, we'll say the recording

comprises stereo overhead mics and stereo room mics; close mics above and below the snare drum; close-mics inside and outside the bass drum; and individual close-mics for the hi-hat, two rack toms, and a floor tom. My thought process and workflow might go roughly as follows:

First off, I route and pan the split-stereo overhead and room channels to stereo group channels in my DAW so that I can process both sides of these stereo files at once. The overhead mics had been set up at similar distances from the snare drum, but it's still a bit off-center so I decide to try time-shifting one side of the overheads file to phase-match the snare hits. First I roughly line up the two sides of the split-stereo audio file by eye. Then I flip the polarity of one side and tweak the time shift by ear for greatest cancellation, before returning the polarity to normal. The snare is now closer toward the center of the stereo image and remains more solid sounding in mono. Checking the rest of the overheads sound for mono compatibility reveals a loss of low end in the kick drum spill and a little less brightness in the cymbals. The changes aren't serious enough that I think they'll drastically compromise real-world mono playback. Besides, the stereo width presented by the overhead mics is fairly extreme, and in narrowing that down more toward my own personal preference I also inevitably reduce the tonal change when switching to mono listening. The room mics seem to keep their sound surprisingly well in mono, but they are leaning a little to the left in stereo, so I nudge up the level of the right-hand split-stereo file to compensate.

Because the overheads seem pretty well-balanced, I decide to start my balancing with them when the time comes. That said, I'm concerned that the snare drum feels a bit lightweight in the overheads, so I next solo the two snare close mics to hear what they sound like together. The polarity of the bottom mic's waveform is, as expected, inverted compared with that of the top mic, and correcting this immediately fills out the drum's low end—good news. While examining the waveforms, I notice that the bottom mic's waveform is also trailing slighting in the timeline, so I time-shift that to see if I can achieve any further improvement. There is only a small change, but it's better than nothing. I still feel that the snare's low midrange could be more powerful, so I insert a phase rotator plug-in on the under-snare channel and adjust the amount of rotation by ear, listening for an even fuller sound. This gives me another small but useful improvement.

The kick mics get soloed next, so I turn my nearfields up a bit to hear the low end more clearly. The signals are both in polarity with each other, but there's about 7ms of delay between the inside-mic and outside-mic hits. I try phase matching again, but it doesn't sound as good and unduly emphasizes an undamped low mid-frequency drum resonance. I undo that time shift, and experiment instead with a phase rotator. I get a really useful result here, with more positive low frequencies and an end to that pitched resonance. The snare and kick mic pairs are also now routed to their own mono group channels,

because I don't normally pan the individual mics separately. Inserting a high-pass filter into the kick drum group channel only begins to affect the drum's low-end thump at about 42Hz, so I roll the frequency back to 35Hz for safety's sake and then duplicate that setting to all the other drum channels. The final preparatory stage is to open the overhead mics and then fade each of the other four close mics up so that their panning can be matched with what's in the overheads.

Right—time to start balancing! First I mute everything but the overheads group channel and set a level to leave lots of headroom on the mix buss. This production is in an upbeat pop/rock style, so I decide that the snare's probably the most important instrument and add that in next. Fading it in supplements the overheads nicely, adding the snare-drum weight that previously seemed to be missing. Balancing the snare with the cymbals isn't easy, though, because hi-hat spill on the top snare mic is putting the hi-hat out of balance before the snare feels like it's at the right level. I try using more of the under-snare mic, but that detracts from the overall snare timbre, so I just put the snare where I want it and leave the hat too loud for the moment.

The kick drum feels like it's probably the next most important instrument, but when I unmute its group channel, the kick drum's tone becomes rather boxy. Inverting the polarity of its group channel improves things considerably, but repeatedly muting and unmuting the bass drum highlights that its spill is destructively phase canceling with the snare's nice low midrange. I try time shifting both kick drum mics a little while listening for a better snare sound, but by the time the snare is sounding better, the kick drum's low-end power is suffering, and phase rotation of the kick drum's group channel doesn't yield any better outcome. So I go for the best snare sound and try to get a balance for the kick, switching back and forth between my nearfields and Auratone substitute to gain more perspective. Once set, the fader level seems fairly stable most of the time, but occasionally the odd uninspiring downbeat makes me yearn to nudge it up a fraction.

Under normal circumstances, I'd probably turn to the hi-hat mic next, but given the overprominence of that instrument in the balance already, I decide to leave it out completely. That leaves only the toms, which so far feel a bit soft edged and distant compared to the snare. Mixing them in one at a time I check for the most promising polarity of each. For both rack toms, the different polarity settings have different effects on the sound of the ensemble, but it's difficult to decide which I like more. Neither seems to affect the snare unduly, so I choose the settings that seem to give a bit more "snap" to the tom sound. I don't waste time with any more detailed phase matching for tonal reasons, because the tom sound just doesn't feel that important to this particular mix. When it comes to balancing, though, I find I can't push the tom-tom faders up far enough without spill making the cymbal balance overbearing. The floor tom close mic clearly detracts from the snare sound when I first fade it in, but inverting its polarity

remedies this situation and otherwise the spill elements picked up by this mic don't seem to cause real problems with any of the other instruments. The main concern while balancing it is that the toms' sympathetic ringing ends up a bit overpowering when the close mic is at a level suitable for the drum fills. As a final touch, I bring up the room mics, but at the low level I choose for them, their exact polarity/phase relationship with the rest of the mics seems to be much of a muchness, so I leave them as they are in that regard. I switch between my nearfields and Auratone substitute a few more times for more perspective, and then I'm ready to add in the next part.

Realistic Expectations

So what has been achieved up to this point in my cloud-cuckoo mix? Well, I've managed to arrive at a reasonable initial balance with a pretty good snare sound. Neither is perfect, by any means, but I've made a lot of important decisions and the sound is already getting better. That's more than enough progress for the moment—after all, I've only just got started! What's absolutely crucial, though, is that the balancing process has already furnished me with the beginnings of a game plan, because I've paid attention to what the instability of certain faders is telling me. So I now know that the kick's low frequencies are balanced too low, and occasional hits are played too quietly; that spill on the snare and tom mics is balanced too high; and that some undamped resonances on the floor tom are balanced too high as well. Notice that I've been able to diagnose all these ailments simply by thinking in terms of balance. Can I hear what I want to hear? If not, then I know I'm going to need further processing. Naturally, the next question is, what processing do I need? The answer to that one is also to be found right there in the balance, but to unveil it you need to learn how to identify exactly what's making each fader unstable—which is the subject of the rest of this part of the book.

Your first draft mix may be totally up the spout, but don't sweat it. However misguided any of your instinctive mix choices, the cold, hard realities of side-by-side comparison will set you straight at the mix referencing stage.

Clearly, the effects of polarity and phase changes during balancing introduce many very subjective tonal decisions into the process, and no matter how much experience you have of the style you're working in, you'll unavoidably gravitate toward the kinds of sounds you personally prefer. That's perfectly fine—there's room for that in a mix, and it's what makes each engineer's work unique. "After all," says Paul Worley, "the music that really works—the music that really rings the bell and makes people go crazy—is the music that doesn't sound like everything else."[36] Of course, every small-studio user's worry is that his or her own favorite sounds will make the track less well-suited to its target market and thus less competitive. In some cases this fear will probably be justified, but you shouldn't fret about this possibility unduly at the

balancing stage—wait until you have a draft version of the fully mixed sonics that can reasonably be lined up against your mix references. Experience counts for a lot in mixing, and well-honed instincts can quickly create most of a release-ready mix. However, it's important to accept that instincts can be unreliable, so there's no shame in being a little uncertain about everything until you can properly acid-test your work against relevant professional productions. If you're just starting out with mixing, your first draft mix may be totally up the spout, but don't sweat it. However misguided any of your instinctive mix choices, the cold, hard realities of side-by-side comparison will set you straight at the mix referencing stage, and as a result your intuition will be improved next time round.

CUT TO THE CHASE

- The simple act of mixing your musical sections and your instruments in order of importance avoids a lot of the mixing pitfalls most frequently encountered by small-studio users. It also reveals the logic behind many arrangement and mixing decisions and helps you make the best use of limited computer CPU power.
- Different types of instrument recording demand slightly different balancing procedures, but in every case you should set up your mixer so that control is straightforward; optimize polarity and phase relationships; remove any unwanted low end; decide on stereo positioning and width; and set the level.
- Polarity inversion, timing shifts, phase rotation, and all-pass filtering are all useful tools for getting the best tone from multimiked instruments and for turning the spill in multimiked ensemble recordings from a curse into a blessing.
- Unwanted low frequencies can usefully be removed from the majority of tracks in a small-studio mix using high-pass filtering. Avoid resonant filters and those with slopes steeper than about 18dB/octave for this purpose. Judge the filter cut-off frequency within the context of the mix, and be especially careful not to rob drum transients of their weight.
- The way different engineers use the stereo field can vary a great deal and largely comes down to personal preference. However, there are some more objective technical issues to consider when panning, such as mono-compatibility, bass playback efficiency, and the danger of nonsensical musical results if only one side of the stereo mix is heard.
- Ensure that you spend quality time with your faders, taking full advantage of all your monitoring systems and techniques to get the clearest view possible of your initial balance. You shouldn't expect a finished mix at this early stage, but if you sensitize yourself to the stability of different fader settings, then that will give you valuable clues about when and how tracks need processing.

Assignment

- Write down all the sections in your production in rank order. Now list the instruments in each of those sections in rank order too. Consider whether any further multing might be sensible in the light of this ranking.
- Build up a balance for the most important section of your production, starting with the most important instrument and then working through your list in order. Deal with routing, phase, high-pass filtering, and panning considerations for each track before trying to decide on the fader level.
- Make a note of any faders that feel unstable.

www.cambridge-mt.com/ms-ch8.htm

CHAPTER 9

Compressing for a Reason

Few aspects of mix processing seem to confuse more small-studio owners than compression, despite the fact that this mix process is actually closely related to what the faders do—the only thing that really differentiates a fader from a compressor is that the latter can perform balance changes automatically. It's easy to lose sight of this basic principle when faced with all the technicalities of compressor designs and controls, so this chapter deliberately focuses on that issue while looking at the practical concerns of using compressors at mixdown. Technical bits and bobs will be sprinkled in as they become relevant.

9.1 COMPRESSION WITH TWO CONTROLS

A compressor is effectively a fader that you can program so that it wiggles around in real time. The beauty of it at mixdown is that you can combat undesirable signal-level variations that would otherwise prevent an instrument from maintaining its position in the balance. In other words, a compressor can iron out some of those pesky fader instabilities that we met while performing the initial balance in Chapter 8.

To explain how a compressor can do its job, let's take the example of a lead vocal recording where the singer mumbles some of the words. If you set your fader so that the majority of this vocal part is nicely audible in your mix, the lower-level mumbled words will start playing hide and seek. If, on the other hand, you fade the vocal up so that the mumbled syllables come though, then the rest of the vocal will eat Manhattan! The result is that no single fader setting will give you a good balance, because the difference between the vocal recording's highest and lowest signal levels (the dynamic range) is too large.

Compressors provide a solution to this problem by reducing dynamic range. In other words, compression reduces the level differences between the mumbled and unmumbled words in our hypothetical vocalist's performance, making it easier for you to find a static fader setting that works for all of them. The way the processor does this is to turn down (or "compress") the louder signals in a given audio recording so that they match the quieter signals more closely. All the compressor needs to know to pull off this stunt is which signals you

Mixing Secrets for the Small Studio.

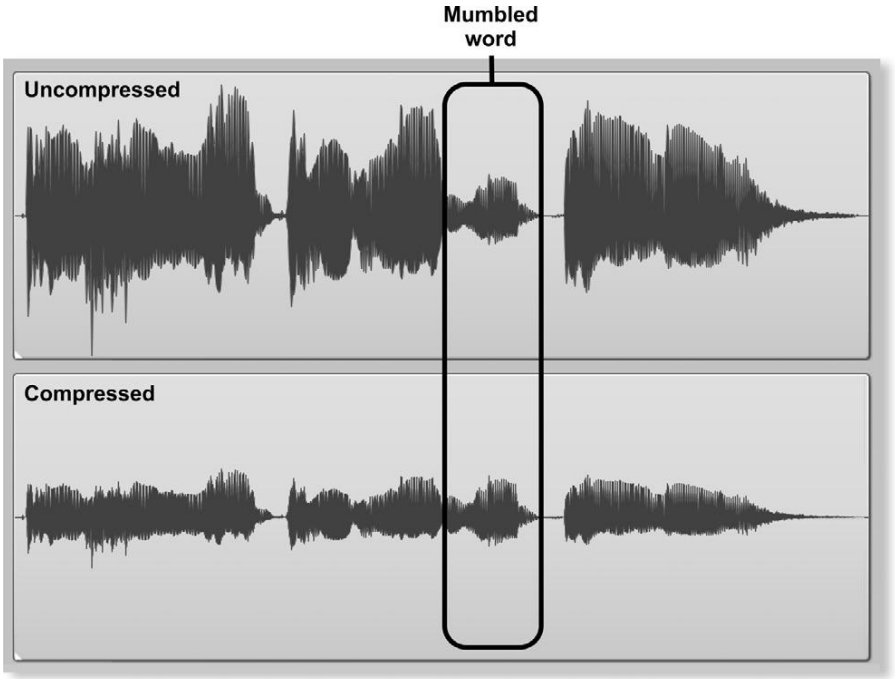

FIGURE 9.1
This example shows how compression can rebalance the relative level of a vocalist's mumbled word by turning down all the other words.

consider to be too loud, and all compressors have a control for specifying this. What can be a little confusing, though, is that this control is implemented on different compressors in quite different ways.

The Many Faces of Threshold and Makeup Gain

There are three different labels you'll commonly see for this first main compressor control:

- *Threshold.* This is the most common control design, but it can seem a bit counterintuitive to start with, because you need to turn the knob down to increase the amount of compression. This is because the control actually specifies the level threshold above which the signal is considered to be too loud. So when the Threshold control is set to maximum, little if anything is considered too loud and precious little compression occurs; as you begin to turn the control down, only the signal peaks are reduced in level; and as the control reaches its minimum setting, all but the softest signals are being sat on.
- *Peak reduction.* A control labeled like this (or sometimes just "Compression") will give you more compression (i.e., more reduction of peak levels) the more you turn it up. Easy peasy!

- *Input gain.* In this case the compressor design has a fixed threshold level above which it will start turning the volume down, so the way you set the amount of compression is simply by varying the input signal's level relative to that threshold. The more you turn up this control, the more the signal exceeds the fixed threshold level, and the more compression you get. Although this sounds a bit like using a Peak Reduction control, the crucial difference is that here the signal's overall level also increases as you add more compression.

I usually recommend that newcomers initially steer clear of compressors with Input Gain controls, because the level hike can easily give the erroneous impression that your processing is improving the sound, even if the amount of compression is way overboard. By contrast, the other two control types tend to make the signal quieter when you compress, so there's less danger of overprocessing. Nevertheless, in the long run it's smart to become comfortable with all three of these common control setups so that you have the widest choice of compressors available to you at mixdown.

It's almost impossible to squish an instrument's dynamic range without simultaneously altering its subjective volume. Most compressors acknowledge this by providing a means of redressing this level change. In most cases, this will simply be a gain control (typically labeled "Makeup Gain" or "Output Gain"), but in some compressors the manufacturers have designed this level-compensation to be automatic. The automatic option seems like a bright idea on the face of it, because it leaves one hand free for beer, but my experience with such designs

FIGURE 9.2
Some examples of compressors using each of the three main control types: Threshold (Waves Renaissance Compressor, *right*); Peak Reduction (Universal Audio LA3A, *bottom*); and Input Gain (Steinberg's Cubase Vintage Compressor, *left*).

Automatic gain makeup seems like a bright idea on the face of it, because it leaves one hand free for beer, but it encourages inexperienced users to overcook the processing.

suggests that they almost always make compressed signals feel louder than uncompressed ones, which again encourages inexperienced users to overcook the processing. Fortunately, many software compressors allow you to choose between manual and automatic gain adjustment here, and I always choose the former if at all possible.

You now understand the theory behind the two most fundamental compressor parameters: threshold and makeup gain. (For the sake of simplicity, I'll refer to them from now on by these names, rather than their variants.) Although many compressors have masses of other controls, you can actually improve your mix balance a tremendous amount without tangling with any of them, so for the moment let's abandon further discussion of technicalities and concentrate on getting the best practical use out of what we already know.

To start with, go back to your initial balance and mute all the channels. Make sure, though, that you leave the fader levels undisturbed at their "best guess so far" settings. Now start reintroducing the instruments in order of importance, but this time concentrate on how compression might be able to improve the solidity of the balance.

Which Tracks Need Compression?

The first thing you'll need to decide for each track you add into the mix is whether it actually needs to be compressed at all, and this is where you can draw on what you've already learned at the initial balancing stage. Remember that the first main aim of mixing is simply to achieve a good balance of the instruments; it's only when you can't find stable fader settings that deliver this goal that you should really consider applying any additional processing. There are lots of possible reasons why a given fader setting can feel unstable, but only some of those problems can be solved with compression, so you need to understand exactly which kind of instability you're looking for in this case.

The simplest clue that you need to compress a track is that you keep wanting to reach over and adjust its fader. So in the vocal example I gave earlier, you might feel you've got the right level for the whole track, but then you find yourself lunging for the fader whenever a mumbled phrase comes along. Because compression can reduce the difference in levels between the clear and mumbled words, it can remedy this problem so that you can leave your fader in peace. To state this idea more generally, compression can be your friend wherever loud moments of a track poke too far out of the balance or wherever quiet details get lost.

If you scrutinize each unstable fader and can say, hand on heart, that none of them suffer from this kind of problem, then no one's going to arrest you for leaving the compression well alone. However, although it's vital that you don't

assume that compression is needed on any track, there are lots of good reasons why compression is one of the most commonly used effects in mainstream record production. In the first instance, no musician ever maintains his or her balance perfectly, and though the top session players can get pretty close to that ideal, I've not heard much evidence that those guys hang round small studios very often. In the real world, you have to do the best you can with whatever performers you have on hand, and that invariably involves a certain amount of airbrushing when you're trying to rival the polished performances on typical commercial releases. The drummer may hit the snare unevenly or may lean on the hi-hat too hard (as in the Chapter 8 case study); the bass player or guitarist may not be able to pick/strum evenly or fret cleanly during a fast riff; and the singer might be lucky to hit the high notes at all, let alone at a consistent level.

The instruments being played can cause just as many problems too. Badly matched components in a drum kit are one of the most frequent culprits, whereas guitars and amplifiers with poorly controlled resonances are also common. Room modes can add further unevenness to the volumes of different instrument pitches. Machine-driven parts aren't immune from problems either, because sloppy MIDI programming can make even the best sampled and synthesized sounds feel lumpy in terms of level. Likewise, there's no end of ways in which a sampler or synthesizer can spoil the balance of a beautifully programmed MIDI part.

But even if there's little remedial work to be done, the demands of many modern music styles are almost impossible to meet without the help of compressors. No bass player I've ever met can play as consistently as is demanded of a mainstream up-tempo pop or rock mix, for example. Lead vocals are seldom mixed without compression either, because singers naturally tend to have a wide dynamic range but most producers choose to narrow this considerably so that the main melody and lyrics remain almost uncannily audible at all times.

FIGURE 9.3
Apple Logic's built-in compressor is one of many software designs that offer automatic gain compensation (via the Auto Gain field in this screenshot). For mix purposes, it's usually better to switch it off.

Getting Down to Processing: First Steps

Once you've located a track that you think might call for compression, then it's time to choose your weapon. So which compressor should you choose? At the risk of uttering studio heresy, I'd say there are more important things to worry about to start with than the particular model of compressor you use. As Tony Visconti puts it, "A compressor is a compressor—it doesn't matter whether it's a $5000 one or a $150 one. You have to know how compression works to use it."[1] To be honest, you might as well use whichever one comes to hand, bearing in mind the advice I offered earlier regarding designs with Input Gain controls or automatic gain compensation. It can make things easier if you use one that has a gain-reduction meter of some type, though, as this will show you when and how hard the compressor is working in turning down the levels of louder signals. Gain-reduction displays are typically in the form of VU-style moving-coil meters or LED bar graphs, and sometimes the compressor's normal level meter can be switched to show gain reduction instead. Whichever one you get, it'll usually be calibrated in decibels and you'll see it move whenever compression's happening to show how much the signal level's being turned down. In case you're still in a quandary about which plug-in to use and you have access to a VST host, then try either Jeroen Breebaart's PC-2 or Tin Brooke Tales's TLS3127LEA as a starting point—they're both freeware and have simple interfaces.

Insert your chosen compressor into the channel in question, and if there are any presets available for it, then select something likely looking—again, there's no need to give it too much thought for now, just go with your gut. When it's

FIGURE 9.4
When you're learning to use compression, choose a fairly simple design to start with, such as these two excellent freeware VST plug-ins: Jeroen Breebaart's PC-2 (*left*) and Tin Brooke Tales TLS3127LEA (*right*).

loaded up, pile on a fair dollop of compression by twisting the Threshold control so that the gain-reduction meter shows at least 6dB of compression occurring on signal peaks, and then adjust the compressor's Makeup Gain control to compensate roughly for any overall level change. Now make use of your different monitoring systems (especially the Auratone-substitute) to help you tweak the fader again, and ask yourself this question: Does the level of that track feel any more stable in the mix?

There are a lot of possible answers, so let's look at each in turn. Clearly, if your compression stabilizes the fader, then you've solved your balance problem and the job's done. However, even if you think this is the case, do still try turning the Threshold control back up a little to see if you can get away with using less compression. Pushing your channel compressors too hard is a common mistake that can slowly suck the life out of a mix if it's duplicated across all your tracks. "Compression is like this drug that you can't get enough of," says Joe Chiccarelli. "You squish things and it feels great and it sounds exciting, but the next day you come back and it's like the morning after and you're saying, 'Oh God, it's too much!'"[2]

Pushing your channel compressors too hard is a common mistake that can slowly suck the life out of a mix if it's duplicated across all your tracks.

If the balance problems aren't yet solved—you still can't stop fidgeting with your fader—then try rolling the threshold down further to see if that makes it easier to find a decent fader level. Feel free to completely max out the control if you like, even it if makes the result sound rather unnatural for the moment. The important thing is to keep concentrating on the balance and whether the compression can deliver the static fader level you're after. Again, if you can get the balance you're happy with and you find any side effects of the compression appealing (as they often can be), then consider yourself the champ and turn your attention back to the rest of the instruments in your mix. In the process of experimenting with the threshold, it may become apparent that different settings actually suit different time segments of the track you're working on, in which case you may wish to do some additional multing of the track to implement this effect.

In the event that you're able to nail down an appropriate balance through heavy compression, you could nonetheless find that the processing isn't doing nice things to the instrument sound in question—perhaps it's making the performance feel lumpy and unmusical or altering the tonality in some unpleasant way. In that case, try switching to a new compressor or compression preset and have another crack. Different compressors and presets can respond very differently for similar settings of our two main compressor controls, and you don't really need to understand why this is to reap the benefits. Compare a few options and choose the one that does the job best. With a little experience, you'll soon build a shortlist of personal favorites for different instruments.

Give it a few months, and you'll be sucking your teeth and saying, "That Fairchild's *phat* on sleighbells!" with the best of them.

When Compression Is Not the Answer

In a lot of cases, though, you won't be able to stabilize the fader fully, no matter which compressor you use or how you set the two main controls. This is the point at which a lot of inexperienced engineers throw in the towel and simply settle for a compromise between dodgy balance and unmusical processing side effects. What you need to realize in this kind of situation is that your mix is trying to tell you that compression isn't the answer, especially if you've already tried a few different compressors or presets. So don't get too hung up on the idea of trying to fix all of an instrument's balance problems with compression, because other types of processing are often required as well. At this stage it's enough just to improve the track's balance in your mix as far as you can without making the sound less likable in a subjective sense.

So let's quickly summarize what we've covered so far. First, concentrate on the balance—can you hear everything you need to hear? When you can't find a static position for a particular fader, then compress (and maybe mult as well) to try to remedy that. When compression does solve your balance problem, then you need to ask yourself a second question: Do I now like the subjective "sound" of my compression? If not, then try a few different compressors or compressor presets. If you still can't find a static fader position that works, then play it safe with the compression until you see what other processing might be able to offer.

9.2 REFINING COMPRESSION SETTINGS

If that were all there was to using compressors, then you'd be forgiven for wondering why manufacturers bother including any other controls at all. Some classic compressors do admittedly only have two knobs (the Teletronix LA2A, for example), but if you've already taken the opportunity to try out a variety of compression presets on the same sound, you'll have noticed that some of them will work more effectively in evening out the levels of the instrument in question than others, and that's because of the deeper compression parameters in each preset. If you can learn to adjust these parameters for yourself, you can match the compressor's action more closely to the specific dynamic-range characteristics of the input signal, and therefore more effectively achieve the static fader level you're looking for. Besides, once you get some practice with the extra controls, it actually ends up being quicker and easier to set them up from scratch anyway.

At low ratio settings (something like 1.5:1) the overshoots are nudged politely back toward the threshold, whereas at higher settings (12:1, for instance), they're beaten back by club-wielding thugs!

Another argument for learning about all the compression controls is that, although the technical *raison d'être* of compression is gain reduction, in the real world compressors usually do more than just reduce gain. They may also change the tone of processed signals quite a lot, even when compressing comparatively little. So if you get into a situation where you like the general tone or attitude that a certain compressor is imparting to your track, but there are no presets that are suitable for the instrument you're processing (or indeed no presets at all), then it's useful to be able to tweak the gain-reduction action manually to suit. That way you can switch between different characterful compressors to find the one which best enhances a given track, while still keeping the balance under control. Let's now introduce some of the more advanced parameters by demonstrating how each allows the compressor to adapt itself to specific tasks.

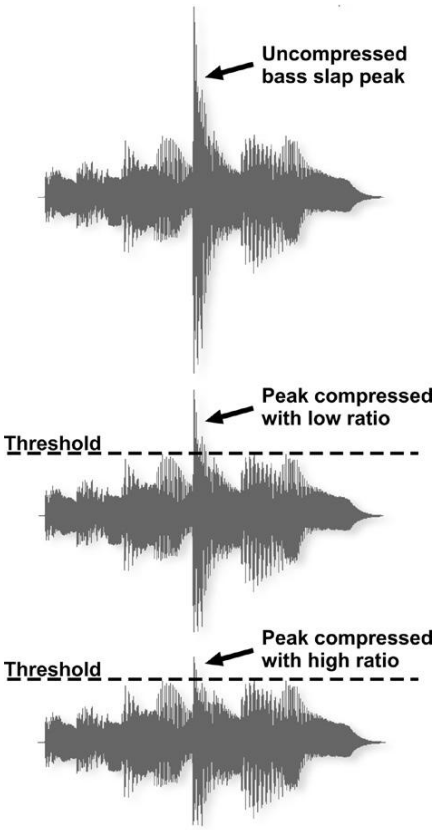

FIGURE 9.5
The slap-bass peak in the top waveform isn't reined in sufficiently at a 2:1 ratio (middle waveform) and requires much higher-ratio processing at 20:1 to bring it into line.

Compression Ratio

As a first example, let's consider a slap-bass part. Now, as everybody knows, the best processing for slap bass is that button labeled "Mute," but let's assume for the moment that this option has been ruled out. This particular slap-bass part balances fine with the rest of the track, except that the odd slap note really takes off and leaps forward out of the mix. The essence of the problem is that you only want to turn down these sporadic signal peaks, but you want to turn them down firmly in order to match the levels of the rest of the bass part. What compressors do is reduce the amount a signal level exceeds the compressor's threshold level, so in this case you want your compressor to put up a proper fight and all but stop the input signal from exceeding the threshold level. That way you can set the threshold level just above the level of the majority of the bass part, and it will then kick in at full force only when the overzealous slap notes hit.

By contrast, imagine an acoustic guitar part where there aren't these kinds of dramatic level spikes, but where the overall dynamic range is still militating against finding a static fader level. In this situation, you want your compressor to act more gently on signals overshooting the threshold so that you can set the threshold just above the level of the softest notes and then subtly squeeze the whole dynamic range down to a more manageable size.

It's a compressor's Ratio control (also sometimes labeled "Slope") that allows it to cope with these two contrasting compression requirements, effectively setting how firmly the compressor reins in signals that overshoot the threshold

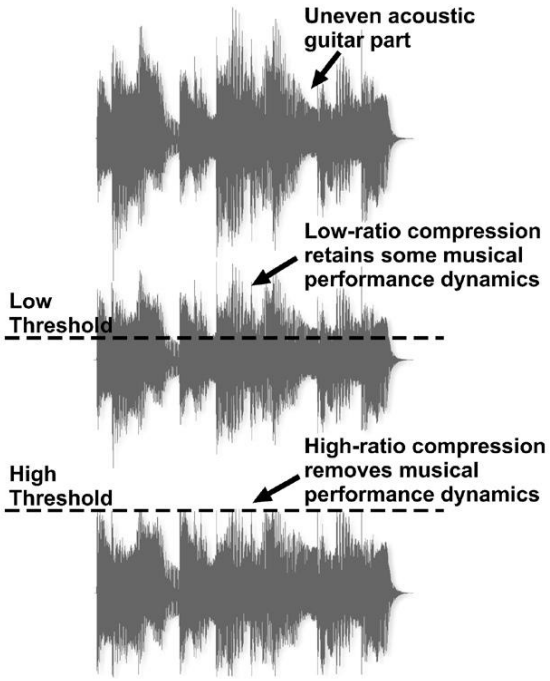

FIGURE 9.6
An acoustic guitar recording (*top waveform*) can be compressed moderately at a 2:1 ratio to add sustain without losing too much of its performance musicality (*middle waveform*). Using a high ratio in this instance (*bottom waveform*) unduly compromises the dynamic nuances of the playing.

level. At low ratio settings (something like 1.5:1), the overshoots are nudged politely back toward the threshold, whereas at higher settings (12:1, for instance), they're beaten back by club-wielding thugs! At the highest settings (some compressors offer infinity:1), overshoots are effectively stopped in their tracks, unable to cross the threshold at all. So for our slap bass example, it'll be high ratios you're looking for, whereas for most routine dynamic-range reduction tasks (such as in the acoustic guitar example), the lower ratios (up to about 3:1) will tend to fix balance problems in a more natural-sounding way.

When I'm talking about a ratio of 3:1 or whatever, what does that figure actually mean? Well, I could draw you some lovely graphs, but frankly I don't think it'd be a tremendous amount of practical help, because some compressors don't label their Ratio controls and different compressors can react quite differently for the same Ratio setting. Instead of thinking in terms of numbers, a more practical and intuitive approach is simply to use a compressor with a gain-reduction meter so that you can see when and how much the compressor is working as you juggle the Threshold and Ratio controls. In the case of our slap bass, you'd start off with the ratio fairly high, and then find a Threshold setting that caused the gain reduction to kick in only on the slap peaks. Once

you'd done this, you'd listen to ascertain whether you'd solved the balance problem, and then adjust the Ratio control accordingly. Still too much slap? Increase the ratio to stamp on the peaks more firmly.

With the guitar example, you might start off with a fairly low ratio (maybe 2:1) and then set up the threshold so that gain reduction was happening for all but the quietest notes. Once the threshold was in roughly the right place, you could then turn back to the Ratio control and tweak it one way or the other to achieve your static fader level. Some quieter notes still too indistinct? Increase the ratio to reduce the dynamic range further and see if this sorts things out. But why not just max out the Ratio control? The danger is that if you turn it up too high, you'll iron out the important performance dynamics that make the part sound musical, leaving it flat-sounding and lifeless, so try to use the lowest ratio that will get the balancing job done.

COMPRESSOR OR LIMITER?

Compressors that are specifically designed to offer very high-ratio compression are often called limiters, so if you find that your particular compressor simply can't muster a high enough ratio to do a particular job, don't be afraid to try a limiter instead. If you do switch to a limiter, though, you'll probably find that it uses an Input Gain control setup, and in some cases the threshold may be set to the digital clipping point for mastering purposes, without any postcompression gain control. This means that you can end up sending your overall signal level hurtling into the stratosphere before you've brought about the gain reduction you require. Fortunately, it's usually easy to add another little gain utility plug-in after the limiter to bring the overall level back down to earth.

FIGURE 9.7
Some good freeware limiters: GVST's GMax (*left*), Tin Brook Tales Pocket Limiter (*right*), and Kjaerhus Audio's Classic Master Limiter (*bottom*).

Compressors in Series

But what would you do if that slap-bass part needed not only high-ratio control of the slapped notes but also more general low-ratio dynamic-range reduction? The answer is that you could deal with the problem by chaining more than one compressor in series. This is common in commercial practice, as it lets you dedicate each specific compressor to a different task. Here's Tom Lord-Alge describing a vocal application: "To make a vocal command attention I'll put it through a Teletronix LA3A and maybe pummel it with 20dB of compression, so the meter is pinned down. If the beginnings of the words then have too much attack, I'll put the vocals through an SSL compressor with a really fast attack, to take off or smooth out the extra attack that the LA3A adds."[3]

A second advantage of a multilayered compression approach is that you can use it to achieve industrial-strength dynamic-range reduction without driving any individual processor too hard. Indeed, many classic recordings have benefited from this principle—a given signal might first have been moderately compressed on the way to the analog tape recorder to maximize signal-to-noise while tracking; the tape itself may then have compressed the signal slightly too, and further layers of compression would have been applied during mixdown. There's another significant reason why some engineers chain compressors: it allows them to blend the sometimes desirable sonic side effects of several different characterful units. This is particularly common with vocal parts, when you're looking for the most successful combination of tone and level control. However, these coloristic uses of compression won't get you anywhere unless you've mastered the processor's basic mix-balancing properties first.

SOFT-KNEE COMPRESSION

Part of what differentiates compressor designs from each other is whether they begin working only the moment the input signal level exceeds the compression threshold (so-called hard-knee designs) or whether in fact they apply small amounts of gain reduction to signals well below the threshold level (soft-knee designs). The main benefit of soft-knee compression is that it makes the onset of gain reduction less obvious and therefore retains slightly more of the natural musical phrasing within the processed part. Quite a few classic studio compressors are by default soft-knee designs, a factor which has contributed to their long-lasting appeal. However, there are many times when the unnatural sound of the compressor working can actually be desirable, so it's not uncommon to find compressors that offer switchable or fully variable hard/soft-knee operation. When using different hard/soft-knee settings in practice, ask yourself the same questions as when trying out different compressors or compression presets: Am I achieving the desired balance, and do I like the processed sound itself?

FIGURE 9.8
A variable Knee control is very useful for adapting a compressor's action to different tasks.

Attack Time and Release Time: Why They Matter

Here's another example for you. Let's say that we're mixing a song where the strummed acoustic guitar has a nice natural sustain that works really well when it's at the right level in the mix, but you find that you have to turn the fader down whenever the player digs in more during your song's choruses. "Sounds like a job for Captain Compressor!" you cry, but when you actually start dialing in the processing you find that, rather than just addressing the level differences between your song sections, our plucky gain-reducing superhero seems intent on evening out the much shorter-term level differences between the attack-transient and sustain portions of each strum. Although you might be able to sort out your overall balance problem with this compressor, you'll be paying an unacceptable price: the impact of each strum will be softened or the instrument's sustain will be unnaturally overemphasized.

The Attack Time and Release Time controls on a compressor provide a remedy to this ailment, because they determine how quickly the compressor's gain reduction reacts to changes in the input signal level: the former specifies how fast the compressor can react in reducing gain, whereas the latter specifies how fast the gain-reduction resets. The reason the compressor in our example isn't doing the job it's being asked to do is because it's reacting too fast to changes in the signal level; in other words, its attack and release times are too small. Increase them and the compressor will begin reacting more slowly, which means that it's likely to deal with this particular balance problem more efficiently, because it will track longer-term level variations (those between our verse and chorus) rather than short-term ones (those between the individual strum transients and the ringing of the guitar strings between them).

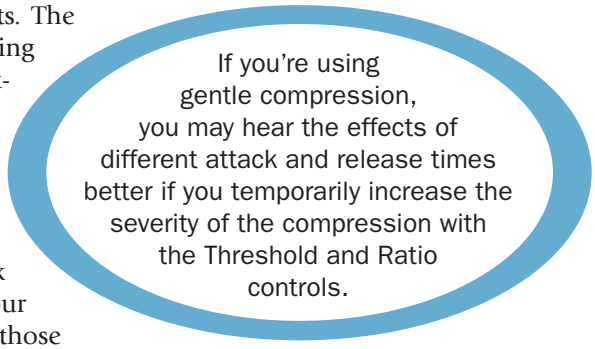

If you look at the scales used on these controls, you may notice that the times are usually expressed in milliseconds, although you do occasionally find microseconds and whole seconds in some cases. However, as with the Ratio control, I wouldn't recommend getting too hung up on exact numbers here, because the figures can only ever be a rough guide to how a specific compressor will actually respond. A much better tactic is to focus on adjusting the controls by ear until you get the best mix balance with the fewest unmusical side effects. If you're using gentle compression, you may find that it helps you hear the effects of different attack and release times better if you temporarily increase the severity of the compression with the Threshold and Ratio controls. A compressor's gain-reduction meter can be a useful visual guide here too, as it'll show you not only how much compression is being applied, but also how fast it's changing in response to the track you're processing.

Drum Compression: Three Different Settings

The ability to adjust a compressor's attack and release times significantly increases the range of balance problems than can usefully be tackled. To illustrate this, let's look at another common example: a snare-drum backbeat:

- *Fast attack, fast release.* If you have a fast attack time, then the compressor will respond quickly to the fleeting initial drum transient, reducing the gain swiftly. If you then set the Release time very fast, the gain reduction will also reset very rapidly, well before the drum sound has finished, such that the lower-level tail of the drum hit won't be compressed as much. Result: Less drum transient.
- *Fast attack, slow release.* Partnering your fast attack with a slower release will cause a rapid compression onset, but the gain-reduction will then reset very little during the drum hit itself, and mostly between the hits. The balance between the transient and sustain portions of the drum will therefore remain pretty much unchanged, and the compressor will primarily just make the level of each drum hit appear more consistent. Result: More consistent performance.
- *Slow attack, slow release.* Increasing the attack time too much will allow some of each drum transient to sneak past the compressor before its gain reduction clamps down properly. This effectively increases the level difference between the transient and the rest of the snare sound. Result: Less drum sustain. (It's worth noting here that, although compression is normally associated with reducing dynamic range, in this case it might easily increase it.)

Here we can see how having variable attack and release times makes possible three different balance results (less transient level, more consistent hit level, and less sustain level), all from the same compressor. This ability to achieve very different effects is partly what confuses some newcomers to compression, and it's also one of the reasons why promisingly named compressor presets don't necessarily do the trick—if your "Rock Snare" preset has been set up to reduce the drum transient, for example, that won't help if you actually want it spikier than a hedgehog's mohawk.

AUTOMATIC ATTACK AND RELEASE TIMES

Many of the most musical-sounding compressor designs treat short-term and long-term level variations differently, adjusting their attack or release times to adapt to the musical material. However, in some cases a dedicated automatic mode is provided which disables the manual attack/release time controls, and this will tend to work best with complex sounds that need transparent gain reduction—perhaps lead vocals or acoustic guitars in a sparse acoustic arrangement. However, if you're after well-controlled transient envelope adjustments (such as those demonstrated via the three contrasting snare-processing examples below), any artificial intelligence can actually hamper you in getting predictable results for large gain changes. "I can't let you do that, Dave."

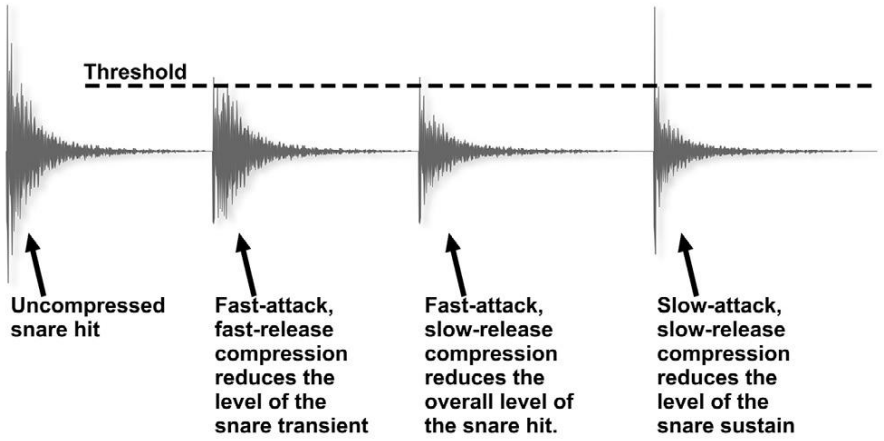

FIGURE 9.9
The effects of three different sets of attack and release times when compressing a snare-drum recording.

Side Effects of Time Settings

Although thinking in terms of balance answers most questions pertaining to attack and release times, in certain circumstances you may find that your settings produce unwanted processing side effects, and it's as well to be on the lookout for these. The first occurs when you set fast enough attack and release times that the compressor actually begins to react to individual waveform cycles, rather than to the signal's overall level contours. The gain reduction then effectively changes the waveform shape and thus produces distortion, the nature of which will depend on the specific sound being processed and the compressor you happen to be abusing. Bass sounds, with their slow-moving waveforms, are particularly prone to this effect, but delicate acoustic instruments can also present difficulties because they will ruthlessly expose the smallest of distortion artifacts.

Another common problem is with percussive bass sounds, such as kick drums, which can appear to lose bass content if you compress them with attack times under about 50ms. This is because the compressor begins clamping down during the first couple of waveform cycles, something that seems to impact on the lower frequencies more than the higher ones, shifting the tonal balance of the sound. Once you know that this is a danger, it's not that tricky to avoid, but if you're not listening for it, then it's easy to overlook while you're concentrating on wider balance issues.

One final observation is that changing a compressor's attack and release times will affect the amount of gain reduction that's occurring for a given combination of threshold and ratio parameters. For example, a side-stick sound (which comprises a short transient and very little sustain) might completely bypass a compressor with a long attack, even if its level shoots way over the compressor's

threshold. So it's not uncommon to keep adjusting compression threshold and ratio alongside attack and release times to take account of these kinds of changes.

LEVEL DETECTION: PEAK VERSUS AVERAGE

Another way you can adjust the time response of a compressor is by altering the level-detector circuit's reaction to short-term signal peaks. If the compressor is aware of every tiny level spike, it'll catch threshold overshoots reliably, but it will also react unmusically to many normal instrument recordings—our own hearing isn't as sensitive to short signal peaks as to slightly longer-term level variations. A lot of compressors, therefore, average their level-detector readings over time to some extent in order to match the response of the ear more accurately and achieve a smoother compression sound, and some designs even provide a Peak/Average (or Peak/RMS) control of some sort so you can vary the amount of averaging manually. If you do have a separate Peak/Average control, though, bear in mind that it will inevitably interact with the Attack Time and Release Time controls, so you may need to hop back and forth between all three to find the best final setting.

9.3 PARALLEL COMPRESSION

Although most studio users first experiment with compressors as insert effects, they can also be used in a send-return loop setup, such that you end up mixing both processed ("wet") and unprocessed ("dry") versions of the same signal. This technique has been in widespread use for decades within the professional community (where it's sometimes called "New York" compression), but it tends still to be underused by small-studio operators. The main advantage of the parallel setup is that the uncompressed signal retains some of the track's original transients and musical dynamics, even if you choose to pummel the living daylights out of the processed channel. "I'm a big fan of parallel compression," says Joe Barresi. "I'll often mix in noncompressed sound with compressed sound to maintain the transients."[4]

As such, one common application of this approach is to help reduce undesirable side effects wherever fast attack and release times prove necessary for balance purposes. Transient-rich instruments such as drums, tuned percussion, piano, and acoustic guitar frequently benefit from parallel processing for this reason, because that way you can use fast compression to bring up the sustain tails between individual hits/notes (a fairly everyday requirement in a lot of styles) without dulling percussive attack or flattening desirable performance nuances. In fact, so good is parallel processing at concealing the side effects of heavy compression that you can frequently afford to compress with gusto in this scenario, all but eliminating transients from the compressed signal so that it can bolster the uncompressed channel's sustain almost entirely independently. The crucial thing to remember is that the sound of the compressed signal on its own here is irrelevant—it's what it sounds like in combination with

the uncompressed signal that counts. "It doesn't matter what it sounds like in solo to me," says Toby Wright. "It can sound like the weirdest thing on earth; as long as it fits and complements the track and what's going on around it, it's there to stay."[5]

Another common reason to employ parallel processing is if you want to overdrive characterful vintage-style compressors in search of juicy tonal and distortion artifacts, because you're free to do so without nearly as much danger of producing a lifeless final sound. In professional circles, this approach most often tends to focus on the most important sonic components of the mix, such as drums and lead vocals, and several parallel compressors might even be blended together. Michael Brauer,[6] Dave Pensado,[7] and Spike Stent[8] all talk about using multiple compressors like this on lead vocals, for example. However, this kind of compressor usage is so subjective that there's little real advice I can offer here beyond Joe Meek's famous adage: if it sounds right, it is right!

COMPRESSING LEAD VOCALS

Although there are no special compression techniques specifically for lead vocals, it's common for their processing to be more in-depth to cope with the difficulties they present. Not only does their naturally wide dynamic range usually need narrowing a great deal, but the listener's everyday experience of natural-sounding vocals makes unnatural side effects from heavy processing difficult to get away with. My main suggestion for getting successful results is to tackle the task in stages:

- Mult sections of the vocal that have very different sounds or levels to different tracks.
- Use an initial layer of gain reduction to even out overall levels—soft-knee 2:1/3:1 compression with moderate attack and release times is likely to work best here if you're after fairly transparent gain control. Alternatively you might wish to use a multicompressor parallel setup to achieve more tone-coloring options.
- If vocal peaks are still poking out unduly, then follow this general compression with a faster, higher-ratio compression setting, or even some kind of limiter. A soft-knee model will be more appropriate for understated level control, but understatement isn't exactly the name of the game in many modern musical styles.

It's as well to realize, however, that compression (and indeed any other type of automatic mix processing) will never be able to achieve the inhuman degree of level consistency demanded of chart-ready lead vocals. The missing piece of the puzzle is detailed fader automation, but it doesn't make sense to get into that until much later in the mix, so we'll leave a proper examination of that subject for Chapter 19. For the moment, just try to do the best you can without incurring inappropriate processing side effects. If you try to nail every last syllable into place with compressors while balancing, you're almost guaranteed to overprocess.

There is one pitfall to be aware of with parallel compression, though: if you fade up either the processed or unprocessed signal, you'll adjust not only the nature of the compression, but also the overall level of that track in the mix.

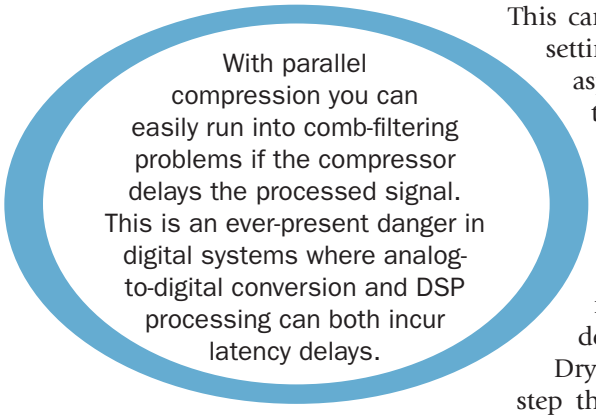

This can make it seem as if even a dodgy compressor setting is helping, simply because fading it up makes aspects of the instrument seem louder. My suggestion for overcoming this effect is to group both the uncompressed and compressed signals to a new channel and then use that group channel to set the balance. The other two faders can then be treated as an extension of the compressor's control set rather than being used for mix balancing. In fact, a number of compressor designs have actually begun to incorporate Wet/Dry Mix controls into their user interfaces to sidestep this issue and allow parallel compression to be employed as an insert effect.

A second thing to be aware of with parallel compression (or indeed when processing the individual tracks of a multimiked recording) is that you can easily run into comb-filtering problems if the compressor delays the processed signal. This is an ever-present danger in digital systems where analog-to-digital conversion and DSP processing both have the potential to incur short latency delays. Fortunately most small-studio DAW users don't run into latency problems too often while mixing because they rely mostly on software plug-in processing and their DAW system will normally have automatic plug-in delay-compensation routines. However, it nonetheless pays to be aware that some plug-ins don't declare their latency to the host software properly, so do keep an ear pricked for the frequency effects of comb-filtering whenever using a parallel processing setup. And even if a plug-in induces no obvious delay, some software nonetheless alters the internal phase relationships between the different frequencies passing through it, so don't be surprised if mixing together compressed and uncompressed channels occasionally alters the combined tone in ways you weren't anticipating.

9.4 BACK TO THE BALANCE

There are clearly a lot of ways that you can refine the way a compressor operates, but none of those refinements will help you at mixdown if you let them distract you from the fundamental questions you need to answer while balancing your mix:

- Is the compression helping the balance?
- Do I like the subjective quality of the compression sound?

If your first compressor choice doesn't give you the balance or the attitude/tone you desire, even once you've tweaked its settings, then switch to another. If you need more clean balancing power or want to use overt compression side effects to your advantage, then try chaining compressors in series or adopting a parallel processing setup. It shouldn't be a huge rigmarole, and there's no need to despair if you can't achieve perfection. The aim is to make the mix better, not perfect,

so don't push your processing too hard at this point. If any balance problem is proving insoluble, then that's probably a sign that simple compression isn't the right tool for the job.

CUT TO THE CHASE

- From a mix perspective, the primary purpose of compression is to achieve a stable balance. If a track balances fine without compression, then don't feel you have to compress it at all. When compressing you need to ask yourself two questions: Is the compression helping the balance, and do I like the subjective sound of the processing?
- The most important controls on a compressor are threshold and makeup gain (or their equivalents), and you should be able to get a long way with compression using just those controls. If they can't achieve a good balance for you without undesirable side effects, then in the first instance just try switching compressor or compressor preset.
- The remaining compression controls allow you to adapt a compressor's gain-reduction characteristics more closely to the needs of the target signal. While this can help you reduce unmusical compression artifacts, it also allows you to take advantage of any desirable sonic side effects of a given compressor in a wider range of situations. Be on the lookout for undesirable distortion or loss of low end when compressing with fast attack and release times.
- There is no rule against using more than one compressor in a single channel. Sometimes different compressors are required to do different balancing jobs, sometimes two compressors can achieve heavy dynamic-range reduction more transparently than one, and sometimes you might want to blend the tonal flavors of two different characterful units.
- Parallel processing can reduce some of the undesirable side effects of fast compression, especially when heavily processing transient-rich signals such as drums, tuned percussion, piano, and acoustic guitar. It also allows you to drive individual compressors harder in order to exaggerate their unique tonal personality in the mix.

Assignment

- Track down all the compressors and limiters on your DAW system so that you know what options are available. If you have little choice, then consider supplementing your selection with third-party plug-ins.
- Check that your DAW system has automatic plug-in delay compensation, and that it is activated.
- Mute all the tracks in your mix, and then rebuild the balance as before, but this time experimenting with compression to see if it can clear up any of the fader instabilities you identified while initially balancing.
- Make a note of any faders that still feel unstable.

www.cambridge-mt.com/ms-ch9.htm

This page intentionally left blank

CHAPTER 10

Beyond Compression

Compression isn't the only automatic gain-changing effect that has applications at mixdown—it's one of a whole family called "dynamics processes," or simply "dynamics." Other members of this family seem frequently to be ignored in the small studio, but they can also have their uses. This chapter explains some of these potential applications.

10.1 EXPANSION AND GATING

Whereas compression reduces dynamic range by turning down loud sounds, expansion does the opposite—it increases the dynamic range by turning down quiet sounds. There are a number of common applications for expanders when balancing a mix, including the following:

- Reducing low-level background noise between vocal or instrument phrases
- Reducing hi-hat spill on a snare close mic
- Reducing the level of a drum sample's sustain tail relative to its transient peaks

An expander's interface is no more complicated than a compressor's. As before, there's a Threshold control to tell the expander which parts of the signal need to be processed, although in this case it's signals below the threshold that have gain-reduction applied to them. The Ratio control adjusts how vigorously the expander turns things down once signals dive below the threshold level. A ratio of 1:1 causes no expansion at all, whereas a higher ratio of, say, 4:1 will dramatically reduce the levels of subthreshold signals. Some expanders even offer a ratio of infinity:1, which effectively mutes any signal lurking below threshold, in which case the processing is commonly referred to as gating, and some specialized expansion processors called gates offer only this most extreme ratio.

FIGURE 10.1
A well-specified software gate, Apple Logic's Noise Gate plug-in.

What can confuse some users initially with expanders and gates is that the Release Time control determines

how quickly gain reduction is applied to subthreshold signals, while the Attack Time knob adjusts how fast the gain-reduction resets. However, given that attack transients in the processed track are often what actually cause the gain reduction to reset in these processors, this control legending is actually more intuitive for a lot of people, so it has been adopted as convention.

EXPANSION: PRE- OR POSTCOMPRESSION?

If you want to use expansion on a track that already has compression on it. Then in most circumstances it's sensible to insert the expander before the compressor in the processing chain. If the compressor squashes the dynamic range first, it will reduce the level differences between the quieter sounds you want your expander to attenuate and the louder sounds you want left alone, with the result that finding a successful expansion threshold becomes more difficult—or even impossible. However, there are inevitably exceptions to this guideline. For example, if you're using an expander to emphasize a drum sound's attack, you might set a high threshold level just to catch the peak of each hit, but if there's any variation in the peak levels, then you'll get some inconsistency in the transient enhancement. Putting a compressor or limiter before the expander in this case could actually make the expander's operation more reliable.

Using expansion for balancing purposes involves answering the same two questions we posed before with compression: Does the processing solve a balance problem, and has the subjective quality of the processed signal suffered in the process? If you can't get the required result with the basic four controls, then try a different expander or expander preset. If that doesn't improve matters, then just do the best you can with your expander and make a note to tackle the problem again with the more specialist tools we'll explore in later chapters.

Stop That Chattering!

There is one practical issue you may have to deal with when using either a high-ratio expander or a gate. If you use fast attack and release times, any signal that hovers around the threshold level can trigger bursts of rapid and extreme gain-reduction fluctuations. At best these will simply sound like ungainly stuttering of the audio, and at worst they'll give rise to a particularly unmusical distortion often called "chattering"—it sounds a bit like whizzing snails in a blender! Although the obvious solution would be to back off the ratio or attack/release times, there are situations where that might not be an option—for example, if you wanted to cut off the unwanted tail of a snare drum sample abruptly.

If you gate with fast attack and release times, any signal that hovers around the threshold level can trigger a burst of rapid gain-reduction fluctuations called "chattering"—it sounds a bit like whizzing snails in a blender!

One fix for gate chattering is to have an additional Hold Time control on your expander, which enforces a user-specified delay between a gain-reduction reset (the gate "opening" to let sound through) and the next gain-reduction onset (the gate "closing" to mute the signal). Just a few milliseconds of hold time can be enough to deal with most cases of gate chattering. Another option sometimes provided is something called hysteresis, which effectively creates separate threshold levels for gain reduction and gain reset. A Hysteresis control will usually be calibrated in decibels, but don't pay the markings much heed—just leave the control at 0dB unless gate chattering becomes a problem, and then turn it up only just far enough to save those gastropods from their unpleasant demise.

Parallel Processing and Range Control

The same kind of parallel processing setup introduced in Section 9.3 with relation to compressors can also extend the possibilities of expansion and gating. Let's say you have a snare drum close-mic recording that catches some important elements of the hi-hat sound between drum hits, but at an unacceptably high level—a fairly common eventuality. You might tackle this problem with expansion, setting the threshold so that the gain-reduction resets with every snare hit and then adjusting the ratio to reduce the hi-hat spill to taste. However, this approach will also increase the hi-hat spill's dynamic range, and that might cause a couple of problems. First, the hi-hat might be picked up with a different tone on the snare and hi-hat microphones (or there may be considerable comb-filtering effects between them), so the louder hi-hat hits, which have more spill, will therefore have a different sound. Second, if the snare and hi-hat close-mics are panned to different locations, the hi-hat image may appear to wander between them as the expansion varies the spill levels.

A solution to this problem might be to set up a high-ratio expander or gate as a parallel process instead, completely removing high-hat spill from the processed channel and then mixing in the resulting "snare-only" signal alongside the unprocessed snare close-mic. By balancing the levels of the processed and unprocessed tracks together, you can freely decide how much hi-hat spill you include in the mix, without altering its dynamic range. Another useful application of this parallel expansion is emphasizing drum transients. Set the expander/gate with a high threshold and fast attack and release times to isolate just a tiny "blip" at the start of each hit, and then mix that back in with the unprocessed sound. Although on its own the isolated transient may sound uglier than the portrait in Lemmy's attic, that's immaterial as long as it sounds good in combination with the unprocessed track.

You'll remember that some compressors have a Wet/Dry Mix control to achieve parallel processing within a single insert effect, and you'll find the odd expander/gate with one of these too. However, it's much more likely that you'll find a Range control instead, which achieves a broadly similar effect by setting the maximum allowable gain reduction. In our hi-hat spill example, then, you'd first set up the expander/gate using the maximum range setting to completely eliminate the hi-hat spill, and once you'd managed to get it triggering reliably you'd roll back the Range control to achieve the desired level of hi-hat spill.

USING GATING TO RESHAPE DRUM ENVELOPES

Almost all Hold Time controls on expanders and gates have a much wider range of delay settings than is necessary to avoid gating chatter, and this opens up a technique that is particularly powerful for remodeling the level envelopes of individual drum sounds. For example, if you put a high-ratio expander or gate onto a snare drum track and set the threshold as high as possible (while still ensuring that a gain-reset triggers reliably for each hit), then you can subsequently use the expander/gate's Attack, Hold, and Release Time controls to dramatically adjust the level envelope of each hit. You could increase the attack time to remove the initial transient completely, decrease the hold time to shorten the drum's tail, and adjust the release time either for an abrupt drum-machine-style ending to each hit or for a more gradual and natural-sounding decay.

10.2 TRANSIENT ENHANCERS

We already have at our disposal various means to emphasize a drum sound's attack transient: slow-attack, fast-release compression (as described in Section 9.2); fast high-threshold expansion; and parallel or limited-range gating. However, there is often a need for specialist tools to adjust the nature of transients in a mix, not least because digital recording doesn't smooth signal spikes in the way classic analog machines used to. "You spend a lot of time carving in analog to make things poke out," comments Jimmy Douglass, "but in digital everything is poking out and sticking in your face already, so the challenge is to smooth it out and stick it back there."[1] Al Stone complains about the impact of digital recording on drum sounds in particular: "If you get the best signal down on multitrack, you end up balancing your multitrack at the bottom of the faders on the way back, because all the transients remain."[2]

Threshold-dependent transient enhancement only works on transients that are higher in level than the rest of the signal, so adjusting complex transient-rich tracks like acoustic guitar or piano can therefore be unrewarding.

Threshold-Dependent Processing

A number of transient-modification options rely on some kind of level-threshold system to work. There is, for example, a processor called an upward expander (or decompressor) that increases the dynamic range of signals above the threshold level, rather than below it. If you set attack and release times very short, you can adjust the Threshold and Ratio controls to boost just the signal peaks. However, if the peak levels aren't absolutely consistent, upward expansion can exacerbate this irregularity, which is often undesirable.

Another type of processor design uses its threshold level simply to detect transients, and whenever one arrives it triggers a momentary gain boost to emphasize it. You usually get control over the amount of "attack" gain applied, and sometimes also the length and shape of the gain envelope used. Although you could create an identical effect using parallel gating, plug-ins that use triggered gain envelopes often have another trick up their sleeves: the ability to trigger a second independent "sustain" envelope when the input signal level ducks back below the threshold. This isn't something you'd use every day, but it has the capability to increase the sustain of drum samples or repetitive synth parts in a much more dramatic and controllable way than is possible with compression.

FIGURE 10.2
Logic's Envelope Shaper plug-in is an example of a threshold-dependent transient processor.

Look Mum, No Threshold Control!

Threshold-dependent transient enhancement has an important limitation, though: it only works on transients that are higher in level than the rest of the signal, so adjusting complex transient-rich tracks like acoustic guitar or piano can therefore be unrewarding, for example. For this reason, a dynamics processor that detects transients in a different way is typically a better bet for balancing them at mixdown. Whether a transient is loud or quiet, it will always by nature involve a rapidly rising signal level, so if a processor looks for this level change, it can detect transients more reliably without the need for any Threshold control. It's then a comparatively simple matter to allow the user to boost (or indeed cut) each transient. (An extension of this idea is for the processor also to detect rapidly falling signal levels, so that you can adjust post-transient sustain levels in the same threshold-independent way.)

LOOKAHEAD AND LATENCY

There are some times with dynamics processing when it's nice to be able to predict the future. For example, if you're using high-threshold gating on a percussive sound, even with a very fast attack time, you can lose some of the sound's initial transient while the gate is still trying to open. To get around this problem, it's handy to get the expander/gate to react slightly before the signal level exceeds the threshold, and some digital processors do provide this option in the form of a Lookahead control. This is usually calibrated in milliseconds, and you only normally need a few milliseconds of lookahead to deal with this kind of problem. Mastering-style limiters commonly have lookahead facilities too, so that they can handle signal peaks more smoothly. While lookahead is undoubtedly useful, however, there is one thing to be careful of: some plug-ins that offer the facility can throw a spanner in the works of your DAW's plug-in delay compensation, so keep your ears open for comb-filtering problems if you're processing in a parallel configuration.

FIGURE 10.3
There are now lots of threshold-independent transient processors on the market, including (*clockwise from left*) SPL's software Transient Designer, Flux's Bitter Sweet II, Stillwell Audio's Transient Monster, and Voxengo's TransGainer.

Jason Goldstein describes one of his favorite applications for this kind of processing: "[I use it] a lot on acoustic guitars. In cases where they're just playing rhythm and they are too 'plucky,' I can… take some of the attack off, without using compression…. It's also a lot to do with the issue of apparent loudness. If you make the attack harder, something will sound louder. It will cut through the mix without having to add additional volume."[3] Jack Joseph Puig used a similar effect on Fergie's "Big Girls Don't Cry": "If you were to listen to the acoustic guitar without the plug-in, it would sound kind of lazy and not urgent…. Instead I wanted to give the guitar an immediate attacking, in-your-face sound, as if it's really digging in and played with fingernails as opposed to the skin of the finger…. The brain subconsciously analyses where the front of the note is and what the feel of a record is. If you have attack like this, the recording feels exciting and vibrant, like it's moving and is fast. It's a large part of the way the record feels."[4]

Such processing is a very powerful addition to the dynamics arsenal, and it is often better suited to rebalancing attack and sustain components within individual notes than things like compression and expansion. You may well not have this type of plug-in built into your DAW, but there is now a good selection of affordable third-party plug-ins available—Stillwell Audio's Transient Monster and Voxengo's TransGainer are two very affordable favorites of mine. It is worth noting, though, that threshold-independent transient processors from different manufacturers seem to respond differently in practice, in terms of exactly how many transients they detect and how they implement the triggered gain

changes, so if you have access to more than one such processor on your system, then compare them to determine which design works best in each mix scenario.

10.3 TEMPO-DRIVEN BALANCING

Here's a final useful dynamics trick: triggering gain changes in sync with the song's tempo. So let's say you have a rhythm-guitar part playing eighth notes, and you want to increase the level of the off-beats in the balance. You try doing this with compression, but that messes with the note sustains. What you really need is a way to switch the track's gain between two different levels in a regular pattern. Fortunately, modern DAW systems now provide several fairly simple ways of achieving this as long as the song you're mixing lines up with the software's metric grid.

One method we've already covered is multing. Slice up the offending guitar part and chuck the bits onto two separate tracks so that you can balance them at independent levels. You might need to fade the edited sections in and out carefully to smooth the level transitions, but the results can nonetheless be very effective and there's the added advantage that you can easily create less regular tempo-related alternations, should you have an unusual fondness for capes and dry ice.

Another option is to find a dynamics plug-in that will respond to incoming MIDI notes—some plug-in gates work like this, for example. Once you have this kind of processor, it becomes easy to set up tempo-matched dynamic effects by programming simple metronomic MIDI parts within your DAW software. In our rhythm-guitar example, you might insert a MIDI-driven gate and then trigger that from an off-beat pattern of MIDI notes to let only the off-beat guitar chords through. You could then back off the plug-in's Range control to reinstate the on-the-beat guitar chords at a more appropriate level.

A dedicated tempo-synchronized gating, chopping, or tremolo effect can also give you this kind of dynamics control, and if you get hold of one that is well-specified, it can also add a lot of extra flexibility by virtue of providing a choice of different level-modulation waveforms. In my experience it's a good idea to

FIGURE 10.4
Programmable tremolo units (such as U-he's Uhbik-T, *left*) and tempo-synchronizable gates (such as MOTU's Pattern Gate, *right*) offer time-based balancing options that extend beyond the capabilities of normal level-based dynamics processors.

try to find a plug-in that also provides a Phase control, so that you can slide the selected waveform shape relative to the beat. This allows you to adjust the tremolo gain changes to suit any instruments that play slightly out of sync with the DAW's metric grid.

CUT TO THE CHASE

- If insufficient dynamic range is stopping you from achieving a good balance for a given track in your mix, then expansion may be the answer. An expander's concept and controls are similar to that of a compressor, so it shouldn't present much of an additional learning curve. In most situations, use expansion before compression in any chain of effects. If you're using a high-ratio expander or gate with fast time settings, then you may need to use Hold or Hysteresis controls to avoid unwanted distortion or level stuttering. Parallel or limited-range gating may be preferable to moderate expansion in some cases. A gate that has a lookahead facility can improve your results when expanding percussive sounds, but if you're using parallel processing, then make sure that the function doesn't induce latency and cause comb filtering.
- Threshold-dependent transient processors have some uses, but you'll get much more mileage from threshold-independent designs at mixdown. Different brands of transient processor tend to sound quite different, so make a point of comparing any alternatives to find the best match for each specific balancing task.
- If you can match up your DAW's metric grid to the tempo of the music you're mixing, then there are various ways you can use tempo-driven dynamics processing to your advantage: via detailed multing, by using a MIDI-driven dynamics plug-in, or by applying a dedicated tempo-synchronized effect.

Assignment

- If you don't have at least one well-specified expander/gate and one threshold-independent transient enhancer in your DAW system, then find third-party plug-ins to fill these gaps in your processing lineup.
- Mute all the tracks in your mix and rebuild the balance again, this time trying out expansion/gating, transient processing, and tempo-driven dynamics in search of a more stable balance.
- Make a note of any faders that still feel unstable.

www.cambridge-mt.com/ms-ch10.htm

CHAPTER 11

Equalizing for a Reason

Speaking in the broadest of terms, Chapters 9 and 10 are concerned with situations in which different time segments of a given track in your mix feel as if they demand different fader settings. It's these time-domain problems that create the kind of fader instability that makes you want to readjust the level setting the whole time, because a fader position that works one moment doesn't work the next. Dynamics processing provides the tools to balance all the audio events in an individual track's timeline relative to each other, which makes it possible to find a single, stable fader setting for that mixer channel.

However, dealing with time-domain issues is only half the story, because there will normally be frequency-domain balance problems to address as well—in other words, situations where an instrument's different frequency regions feel as if they demand different fader settings. The fader instability in this case has a different subjective character to it: you'll struggle to get a satisfactory level setting anywhere in the timeline, because a fader position that works for one frequency region doesn't work for another. This kind of problem requires a whole different set of mix-processing tools, the most common of which is equalization (or EQ for short). Because equalization is so important for mixing, and so many small-studio users seem to get the wrong end of the stick with it, I'm going to dedicate the entirety of Chapter 11 to the subject, before digging into some more left-field frequency-domain tricks in Chapter 12.

11.1 FREQUENCY MASKING AND BALANCE

An equalizer adjusts the levels of different frequency regions relative to each other. As such, it can tackle mixdown frequency-domain imbalances head on, so it's hardly surprising that it's used so much. However, most small-studio operators seem habitually to overuse it, largely because of a fundamental misconception about what it's there for. The root of the problem is that inexperienced engineers primarily use mixdown EQ to try to improve the tone, richness, and general subjective appeal of each of the instruments in their production. Although this is a worthy goal, it's actually a secondary consideration, because EQ has a much more crucial task at mixdown: achieving a stable balance.

Mixing Secrets for the Small Studio.

To understand why equalization is so vital to creating a good mix balance, you need to know about a psychological phenomenon called "frequency masking," which affects our perception whenever we hear several instruments playing together at once. To put it simply, if one instrument in your mix has lots of energy in a certain frequency region, then your perception will be desensitized to that frequency region of the other instruments. In other words, if you have cymbals thrashing away and filling up the frequency spectrum above 5kHz, you'll perceive this frequency range a lot less well in the lead vocal part—the cymbals will be "masking" the vocal above 5kHz. Although the vocal might sound lovely and bright on its own, the moment the cymbals are added it will appear dull. To retain apparently the same vocal sound against the cymbals, you must either reduce the levels of the cymbal frequencies above 5kHz or exaggerate those frequencies in the vocal sound.

EQ presets are of no use whatsoever at mixdown, because the designer of the preset can't possibly predict how masking will affect any given sound in your specific situation.

The ramifications of frequency masking for mixing are enormous. First, it should be clear that EQ presets are of no use whatsoever at mixdown, because the designer of the preset can't possibly predict how masking will affect any given sound in your specific situation—you might just as well ask your Aunt Mavis to set up the controls so they look pretty. By the same token, an EQ setting that worked on one mix can't be expected to work on the next. Mick Guzauski says, "You shouldn't say, 'Oh, I'm gonna add 5kHz to the vocal and pull out 200Hz, because that's what I did on the last song.' It may not work. Your approach has to be modified somewhat for each song, as each will have different instrumentation, a different singer, a different key, and so on."[1]

The effects of masking also mean that even if each individual instrument in your arrangement sounds good enough to eat on its own, you'll still need some equalization to compensate for frequency masking between the instruments in order to maintain the apparent tone of each one within the final mix. What's more, carrying out this necessary processing may make individual sounds a lot less subjectively appealing when soloed, either because certain frequency regions have been exaggerated to overcome masking from less important instruments or because some frequencies have been cut to avoid masking more important instruments. In short, a good mix EQ setting is not necessarily the one that makes the instrument sound best in its own right. In some cases the only way to fit a subsidiary instrument into a crowded mix is if its frequency balance is mangled to the point of absurdity. That's why trying to make every individual instrument sound fantastic is a fool's errand.

Says John Leckie, "Sometimes the drums sound spectacular, but it's not what you want, because you want the guitar to be spectacular, and you can't have spectacular everything—then you wonder why the mix doesn't sound any good, because everything's crowding everything else. When you solo the instruments,

everything sounds good, but when it's all put together it's a jumbled-up mess, so something's got to give way."[2] "Or the opposite can happen," says Tchad Blake. "You solo the kick drum and it'll be just awful. But then listen to it with the bass and it can be fantastic."[3]

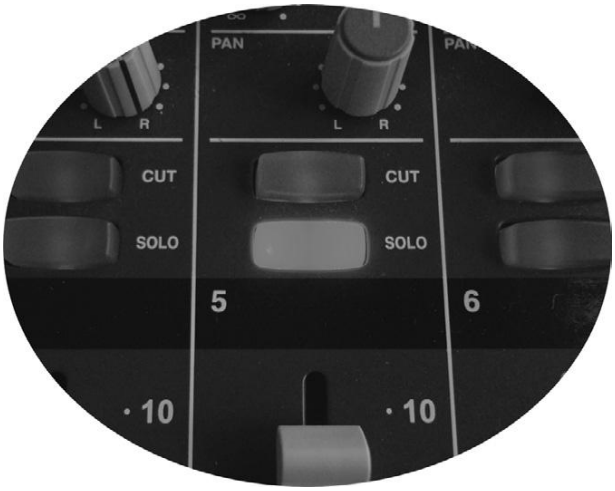

FIGURE 11.1
Don't try to evaluate the success of your mix EQ with the solo button down.

Now I'm not saying that EQ has no role to play in improving the subjective timbre, character, and general splendidness of individual instruments, but this needs to remain well below the balancing role in your list of priorities, otherwise the mix as a whole will never live up to its potential. Ideally, the tone of each instrument should have been optimized at the tracking stage, so if you start mixing one of your own projects and discover a track that's sonically underwhelming (even when it's soloed), then it's often a good indicator that the production's not actually ready to be mixed. Clearly, though, life in the small studio is often far from ideal, and there are usually a few instruments that need extra sparkle, warmth, aggression, moistness, flocculence, or whatever—but EQ often isn't the best tool for adding these qualities. So this chapter is going to concentrate primarily on the balancing role of EQ, and I'll leave it to Chapter 12 to describe some alternative tone-sculpting methods that can deliver serious subjective tonal changes on those occasions where you simply don't like one of the raw sounds you're presented with.

Another lesson to draw from frequency masking is that there's little to be gained by making EQ judgments with the solo button down, because you can't judge the impact of frequency masking between any two instruments unless you listen to them at the same time. Soloing tracks while EQ'ing can be useful for hearing exactly how and where you're altering the frequency balance, but it's imperative that you subsequently validate and refine those EQ decisions within the context of the mix. Finally, frequency masking provides another good justification for building up your mix balance by introducing the tracks in order of importance. If each newly added instrument is less important than the previous one, then you can be pretty sure which one needs to capitulate when masking conflicts arise.

11.2 BASIC MIX EQ TOOLS AND TECHNIQUES

With the key balancing role of EQ firmly in our minds, let's go back to the mix balance we've already fortified with dynamics processing. Once again, mute all the channels and rebuild the balance in order of importance. As you introduce each new track, listen carefully in order to answer these two questions:

- Can I find a fader level for the new track that allows me to hear all of its frequency regions as clearly as I want to?

- Is the new instrument leaving the perceived frequency balance of the more important tracks essentially unscathed?

If you can reply in the affirmative both times, then back away from the EQ with your hands in the air. A surprisingly large number of tracks in most productions only require a bit of high-pass filtering to mix, and you should already have attended to that troubleshooting task while you were putting together your first balance. When your answer to either of the questions is no, however, then it's a clue that EQ might be worth a try.

Shelving Filter Basics

An equalizer is effectively a kind of processing buffet, offering a selection of different filter types for users to dish out as they like. Each of these filters can change the frequency balance of a sound in a different way, so the first decision you have to make with any equalizer is which filter types to use. The best one to start off with when balancing a mix is called a "shelf" or "shelving" filter, a broad-brush processor that can change the level of one whole end of the frequency spectrum. The filter comes in two variants: a low shelf, which affects the lower end of the spectrum, and a high shelf, which affects the higher end. Shelves have at least one user-variable parameter, gain, which simply determines the level change applied to the filtered frequency region. The other main attribute of the filter is its "corner frequency," the notional boundary between the processed and unprocessed frequency ranges. Some EQ designs have shelving filters with fixed corner frequencies, but this usually restricts your options too much for critical balancing work, so seek out an equalizer that has a freely variable Frequency control instead.

Although a few equalizers take the "processing buffet" concept into all-you-can-eat territory by letting you add new filters indefinitely, most equalizers restrict you to a finite number of simultaneous "bands," each of which can contain one filter. It's not uncommon for the bands in real-world EQs also to restrict your choice of filter type and to limit the frequency range over which that filter can operate. In such cases, you may find that shelving filters are only available for some of the EQ bands or that other filter types are selected by default. However, it's rare for an equalizer to have no shelving filters at all, so you should be able to dig one up if you look hard enough.

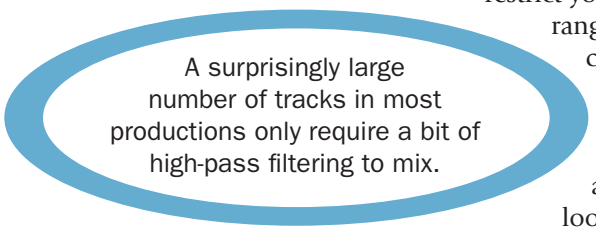

Balancing with Shelves

Once you've got suitable shelving filters on hand, return to that unstable fader and pull it right down. Slowly bring it back up into the mix again until you get to a level where it feels as if some frequency range of that instrument is too prominent in the mix. Take a guess at which half of the frequency spectrum is the culprit ("Is it treble or bass?"), and then load in an appropriate shelving

filter to try to get some control over it. (If you're dealing with a stereo signal, then make sure to insert the processor in such a way that both sides of the stereo file are processed identically; otherwise the stereo image may become less solid.) First wiggle the filter's Gain control around a bit so you can hear what frequencies it's operating on, and then tweak its Frequency knob so that you get some useful handle on the region that's out of kilter, while at the same time minimizing changes to the rest of the spectrum.

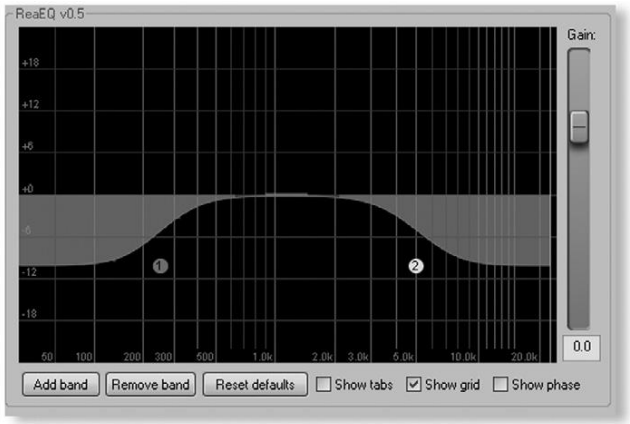

FIGURE 11.2
An EQ curve showing both a low shelving filter (*left-hand side*) and a high shelving filter (*right-hand side*).

It's normal for you to lose a certain amount of perspective during the frequently odd-sounding process of finding a good corner frequency for your shelving filter, so once you've set it, return the filter's gain to zero and give your ears 10 seconds or so to reattune themselves to the actual scale of the original balance problem. When you're back in the zone, gradually start lowering the filter's Gain control to pull the overprominent frequency region down to a more suitable level in the balance. Now turn up your fader a bit more; if that same region pops out of the balance unduly again, then pull back the filter's Gain control some more. Keep turning the fader up, and the filter Gain control down, either until you achieve a good balance for that track or until some other balance problem becomes more pressing.

If you do achieve a good balance straight away, then you need to make sure you're being as objective as possible about the decision. You should switch off the EQ now that you've finished adjusting it, and once more give your ear a few seconds to acclimatize to what the original balance problem sounded like. Then spend a few seconds trying to imagine what the fixed balance will sound like before switching the EQ back on. What you'll find surprisingly often is that your ear lost some perspective while you were fiddling with the filter gain, so your balance fix turns out to be lamer than a one-legged Vanilla Ice impersonator. It can be a bit dispiriting, I know, but your mix (and indeed your mixing skills) will progress more quickly if you force yourself to swallow the bitter pill of fine-tuning your EQ decisions properly in this way. Checking the validity of the balance on your different monitoring systems will also increase your objectivity.

Should a great balance remain elusive, then decide whether the shelving filter is actually helping at all. If the filter is improving the balance in some way, then you might as well leave it in place, albeit subject to the same objectivity checks mentioned in the previous paragraph. If the filter isn't helping you achieve a more stable fader setting, then don't think twice about abandoning it. Remember that shelving filters go on working all the way to the extremes of the audible spectrum, so listen particularly for any undesirable loss of low-end weight or any snide remarks from the studio bat. Whichever of these outcomes

you're faced with, though, you've still got an unstable fader, so you're going to need further processing to deal with that.

In the event that restraining one protruding frequency region with EQ reveals another different region poking out too far from the mix, then by all means have a go with additional shelving filters to see whether you can solidify the balance further. As long as you concentrate on making sure you're actually improving the balance, there's no reason not to have several shelving filters active at once on the same track—the gain changes of each new filter will simply accumulate as you'd expect. Normally, though, there's a limit to how much you can do with wideband processing like this, in which case the more surgical action of a peaking filter may be required.

Adding in Peaking Filters

The peaking filter is a more flexible tool for dealing with localized frequency-domain problems. It creates a peak or trough in the frequency response

EQUALIZING BASS INSTRUMENTS

Judging by the problems I most commonly hear in homebrew mixes, bass instruments typically present the biggest EQ challenges for small-studio occupants. One of the big misconceptions is that all the important EQ adjustments for bass instruments are at the low end. "I also add plenty of top end, so it fits in the track," says Chris Lord-Alge. "You may think that the bass sounds bright when you solo it, but once you put in the heavy guitars it always seems dull all of a sudden."[4] Rich Costey stresses this point: "It's sometimes quite shocking to realize how much top end you need to add to bass to make sure it cuts through a track. If you then hear the bass sound in isolation it may sound pretty uncomfortable, but in the midst of the swirling din of a dense track that amount of top end usually works fine."[5] The additional advantage of defining bass instruments using their upper spectrum is that they will come through much better on small speaker systems, which helps the bass line reach the widest audience. Similar principles also apply to kick drums, of course.

The other big EQ task for bass instruments in most mixdowns involves fitting the kick drum and bass together. "Things can become very tricky with the bass," says Marcella Araica. "It's very easy to lose your kick drum when you add the bass. If things are not EQ'd right, you get more of the bass than the actual thump of the kick. You need to treat the bass and the drums completely differently, otherwise you'll get into trouble."[6] Steve Hodge gives a few more specifics: "Listen to the kick and the bass together and adjust the equalization on the kick drum so that you create a hole for the bass to fit in. I usually cut the EQ around 400Hz just… a couple of dBs, so that it doesn't begin to sound hollow, and see how the bass then fits against the kick. It's usually the perfect fit." He also suggests that adjusting instrument pitches may lend a helping hand here: "The kick has its own pitch, just as the bass instrument does. This becomes more pronounced with drum machines…. Get the bass and the kick into the same pitch range and that will solve a lot of subsonic issues."[7]

(depending on the gain applied), and you can change the center frequency of this adjustment region to target spectral problems. Well-specified models also let you control the frequency bandwidth that the filter affects, and some designs allow such precision in this regard that you can actually rebalance individual harmonics within a sound if you so wish. For mixing purposes you should choose an EQ that allows you to access all three important variables of its peaking filters—the first two controls will usually be labeled Gain and Frequency, but the last may go under the name of Bandwidth, Q, or Resonance, depending on the manufacturer.

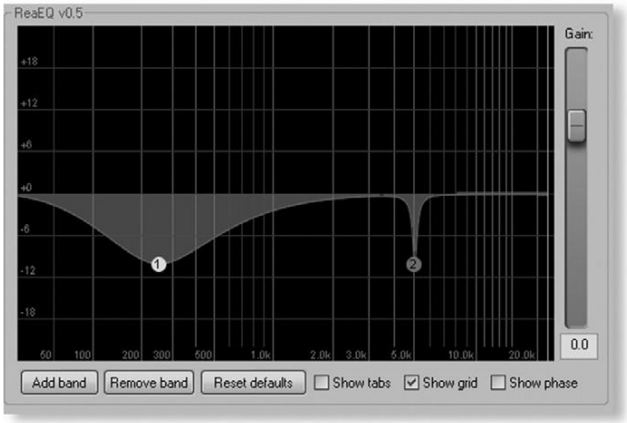

FIGURE 11.3
An EQ curve showing two peaking-filter cuts. The left-hand one has a wide bandwidth (in other words, low resonance or low Q value), whereas the right-hand one has a narrow bandwidth (high resonance or high Q value).

You can follow exactly the same procedure when using peaks for balancing as when using shelves, the only additional consideration being how to set the Bandwidth control:

- Turn up your fader until some frequencies start to poke out more than others.
- Switch on the peaking filter and waggle the Gain control around while you try to find sensible frequency and bandwidth starting points.
- Once the filter is pointing pretty much in the right direction, reset its gain and give yourself a few seconds to become reaccustomed with the original balance problem.
- Turn the peaking filter's gain down (and possibly turn the fader up too) in search of a better balance, and make use of all your monitoring systems to keep perspective.
- When you think you've homed in on a good filter setting, deactivate it for a few seconds and then reengage it to check its validity.
- Refine the filter settings and repeat the previous step if necessary. Remember to keep switching monitors and trying different monitoring volumes.

If possible during this last step, try to keep the bandwidth as wide as you can, while still getting the job done, because undesirable filtering side effects tend to be worse for narrower settings. "I generally use very wide Qs on everything," says Chuck Ainlay. "I don't usually get into narrow-bandwidth boosting or cutting because it takes the music out of most instruments, unless there's some real problem element to an instrument that I need to get out."[8]

Some Applications of Notching

Ainlay's final remark leads us to what is probably the peaking filter's most useful application: making narrow-bandwidth cuts, sometimes referred to as notches. "[I'll often apply] hard notch EQ," says Tom Elmhirst. "I try to find

things that I hate, and the way I do that is by adding specks of EQ with a high Q setting, find what I don't like and off it goes."[9] Notches are regularly a lifesaver on project-studio drum recordings, where there always seem to be a few dissonant drum resonances that stick out of the mix awkwardly even when everything else appears to be at the right level. Because you're usually just dealing with a single frequency, you can position a super-narrow peaking filter precisely on top of it and pull it down out of harm's way without affecting the overall tone of the drum at all. To do this kind of frequency sharp-shooting, though, you'll need a peaking filter that can be made extremely narrow and some means of working out which exact frequency you need. Miss the right setting by a few Hertz, and you may not get the desired effect. One method of finding drum resonances is to boost heavily with your narrowest possible peaking filter, and then hunt for the offending frequency by sweeping the filter around. Although this method certainly works, I've given up using it now because I can find these lingering frequency peaks more quickly and reliably using a high-resolution spectrum analyzer. (My affordable tool of choice here is Schwa's cross-platform Schope plug-in, because you can zoom in with its display to really home in on problem regions of the spectrum. Many other analyzers, including most of the ones built into DAW systems, simply aren't accurate enough for serious notching work.)

GRAPHIC EQ

The type of EQ that's of most use when mixing is called "parametric" EQ, and it is the main subject of this chapter. However, there is another type called "graphic" EQ, which abandons the idea of Gain, Frequency, and Bandwidth controls in favor of a more visually intuitive setup whereby the audible spectrum is split into lots of little slices, the gain of each being controlled using its own tiny fader. With all the faders arrayed in a line, you get a physical representation of the way the frequency response is being processed, which can be read and understood at a glance. This visual immediacy endears graphic EQ to front-of-house engineers, who already have their work cut out trying to stop Florence scaling the lighting rig, but it comes at the sonic expense of coarser control and increased phase/resonance problems, so I'd recommend steering clear of it at mixdown if you can.

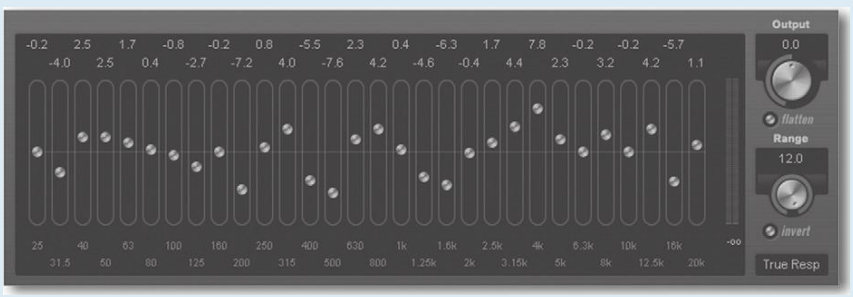

FIGURE 11.4
Steinberg Cubase's built-in 31-band graphic EQ plug-in.

Tampering with the levels of individual frequencies doesn't usually help you when dealing with pitched parts, because a fixed-frequency notch will inevitably affect different harmonics on different notes. Nevertheless, don't rule out notching if the part in question is simple. For example, I quite often find notches handy for bass parts where the fundamental frequencies of different notes need evening out. Any part with a static drone or repeated note within it is also fair game if you want to rebalance the stationary note against any others, although you might find you need to attack more than one of the note's harmonics to achieve the balance you need. Induced noise from the mains electrical supply is another common candidate for multiple notching too, as this kind of noise often comprises a strong hum component at 50Hz (60Hz in the United States) with a series of harmonics at multiples of that frequency.

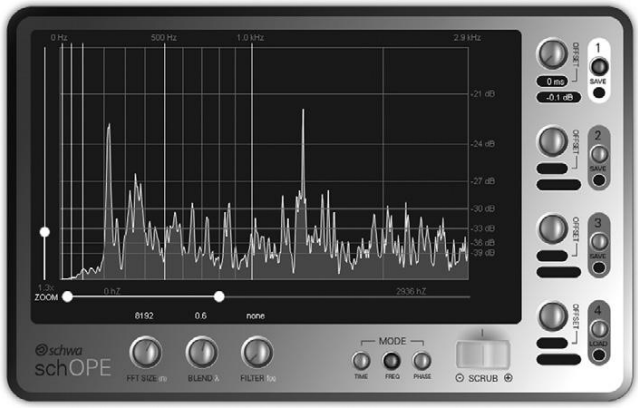

FIGURE 11.5
A high-resolution zoomable spectrum analyzer such as Schwa's Schope is a great help when you're hunting for unwanted narrowband resonances—the snare-drum hit in this screenshot has one at around 1.3kHz, for example.

Narrow-bandwidth peaking cuts are great for counteracting unwanted recorded resonances, such as the nasty honks that room modes can overlay onto a recording that's been carried out in unfavorable surroundings. Guitar amp cabinets sometimes have unattractive resonant lumps in their frequency responses too, which make anything but the simplest part sound uneven. Sometimes singers overplay certain resonances of their vocal apparatus, so that a few single frequency components pop out from time to time and prevent the voice from sitting properly in the mix. (Tom Elmhirst used four different notches to deal with a situation like this while mixing the lead vocal on Amy Winehouse's "Rehab," for instance.[10]) Overall, you just have to be aware that even a single frequency can stymie your balance if it's far enough out of line, so try to make a point of listening for both wideband and narrowband frequency imbalances while working with peaking filters.

Shift Your Listening Perspective

So far we've looked at how to maximize the audibility of each new instrument you add into the mix by ensuring that all its key frequency components are at the right relative levels. Although that's one big part of what mix EQ is about, the other vital step is to check whether the latest track has masked any indispensable aspects of the more important tracks that preceded it. The best way to assess this is to consciously direct your attention to each of the more important parts, listening for undesirable changes as you mute and unmute the latest addition alongside. With some experience you may be able to speed up this process by listening for changes across the board while muting/unmuting the most recently added track, but if the arrangement is complex, then a piecemeal approach will still justify the time required.

WHAT ORDER SHOULD I PUT EQ AND DYNAMICS PROCESSORS?

The main issue to consider when ordering EQ and dynamics plug-ins within a track's processing chain is that your dynamics processors may respond differently if you alter the frequency balance they are fed with. For this reason it makes sense to put EQ last in the chain if you're already happy with the way the dynamics are operating. However, it's not unusual for a frequency imbalance to prevent successful dynamics processing. A cabinet resonance on an electric guitar recording may well cause a single frequency to dominate the timbre only sporadically, depending on which notes happen to hit the resonance. If you compress a recording like this, the compressor won't be able to even out the subjective level of the part, because it will act too strongly on the resonant notes. Dipping the resonance with a narrowband peaking filter precompression would improve this situation.

So the general principle is this: if you're happy with the way your dynamics processors are responding, then EQ after them; if you aren't, then try EQ'ing earlier in the chain to see if you can improve things. In the latter case you may need to reassess your dynamics settings in the light of the altered frequency balance. There's no rule against EQ'ing at multiple locations in the chain either, because you may need one EQ setting to achieve a musical compression sound, but another completely different one to slot the compressed track's frequency spectrum into the mix balance as a whole.

Let's take an example. Maybe you've just added some piano chords into the mix and EQ'd them to stabilize their fader and make them as audible as you want within the balance. However, in doing so you realize that they're masking the upper frequencies of your lead vocal, making it sound dull (or in the case of Morrissey, even duller). What this means is that there's simply not enough space in your mix for your ideal piano tone, and you're going to need to compromise some of its upper frequencies to clear enough room for the lead vocal timbre to shine through. How much you sacrifice the less important track will be a judgment call, of course, depending on which vocal frequencies are really critical and how much they're being masked. You might be willing to sacrifice a small amount of vocal clarity in order to avoid inflicting grievous bodily harm on the piano timbre, and that's a valid decision—after all, most listeners have highly developed expectations of how a piano should sound. On the other hand, you might be able to utterly purée a synthesizer or electric guitar with EQ under similar circumstances, because a listener's expectations of these instruments are usually much less well-defined.

You may be able to utterly purée a synthesizer or electric guitar with EQ, because a listener's expectations of these instruments are usually much less well-defined.

Once you get the idea of shifting your listening perspective so that you're listening to one track while equalizing another, this kind of masking problem should prove no trickier to address than

any of the other EQ'ing we've talked about so far. It is, however, easy to forget to check for masking as you go, and by the time every track is mixed in it can be almost impossible to disentangle complex interwoven masking effects if the mix feels generally cluttered and lead instruments are struggling to sound their best. So try to be as disciplined as you can about this while you build your balance.

Good EQ Habits

You've probably already noticed that my suggested EQ procedure has so far used only EQ cuts, and that's no accident. Although it's actually more natural to think in terms of boosting frequencies you want more of, there are good reasons to retrain yourself to the opposite approach of cutting frequencies you want less of. As I mentioned in Section 4.2, all of us have a strong tendency to like the louder of any two sounds we're comparing, so if EQ boost makes your overall sound louder, then that'll add an element of bias to your judgments. Any EQ boost will also always make the processed track more audible in the balance, so it's trickier to be sure whether you've chosen exactly the right settings. The bigger the boost, the more difficult it is to keep your perspective. Sticking to EQ cuts avoids this bias, or rather biases you against your own EQ curve so that only settings that really work are likely to pass muster when you switch the plug-in on and off. "People need to know how to use their EQ," says Ed Cherny, "learning how to dip some frequencies and move the fader up instead of boosting all the time…. That's the biggest mistake that rookies make—they just reach down and boost the EQ and say, 'That's better!' because it's louder now, but they're not necessarily listening to how the sound is being shaped."[11]

EQ FOR LENGTH

Because different frequencies move at different speeds, changing the frequency balance of an instrument can have an appreciable effect on subjective note durations. As Dave Way explains, "When I think of bass, I think of length; like how long a note should be, because you can make a note longer by boosting 20 or 30Hz. Those frequencies are so big, they ring out longer and make the note actually longer. Sometimes… you might have to clean up the bass and take out some of that 20–30Hz to make it tighter."[12] Eddie Kramer gives a common example where the subjective length of a note may need to be changed in this kind of way: "If the bass drum pattern is busy, you don't want too much bottom end. If the pattern is such that there's a lot of air and space between each note, you can afford to put a little more woof into it."[13]

There's a second substantial reason to avoid EQ boosts: it reduces the damage from a number of troublesome EQ-processing artifacts. Equalizers do a lot of things beyond just adjusting the frequency balance. For a start, most of them adjust the phase relationship between the track's frequency components (often called the track's "phase response"), and it turns out that the biggest

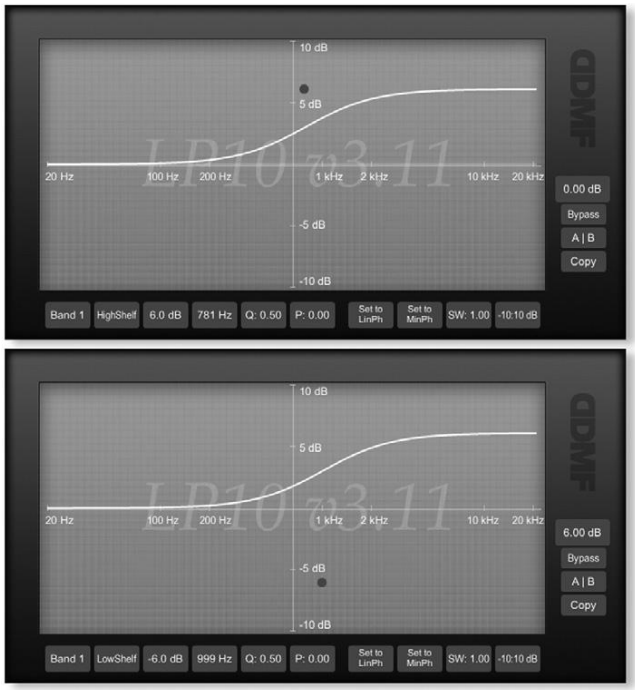

FIGURE 11.6
A lot of EQ curves can be created using either cuts or boosts, as in this example. The upper screen shows a 6dB high shelving boost; the lower one shows the same curve implemented using a 6dB low shelving cut and 6dB of makeup gain, usually the more preferable mixing approach.

phase shifts for shelf and peak filters are applied to the frequency region that's being adjusted. This clearly has knock-on effects for the way a track combines with others in a multimiked recording, because the EQ will alter the sound of any comb filtering between the processed mic and others in the setup. If you boost to give a frequency region in such a track extra importance, not only will the comb filtering at those frequencies alter unpredictably (and potentially for the worse), but you may also increase the severity of the comb filtering in that region by virtue of increasing its relative level. Cutting, on the other hand, concentrates the phase shifts into frequency regions that are less crucial (that's usually why you're cutting them, after all), so any undesirable changes to the nature of the comb filtering won't matter as much and will at any rate be less severe by dint of the level drop.

Phase response doesn't just matter in multimiked recordings, either. Messing with it even on individual overdubs will usually make the sound appear less focused and solid in the mix, and transients in particular can start to become undesirably softened. So here too it's sensible to cut rather than boost, and also to sort out any frequency imbalances with shelving filters first if possible, because peaking filters usually cause more destructive phase shifts than shelves by nature of their design.

But there's more. Filters are also resonant devices, which means that any adjusted frequencies ring on slightly, extending over time—especially if you use narrow bandwidths. Again, this time smearing isn't actually desirable in a lot of cases, so it makes sense to position it in a region of the frequency response where it doesn't impact as much—in other words, by using cuts instead of boosts. In addition, filters can often add distortion to the frequency regions they process, and while this is actually part of the appeal of celebrated analog EQ designs, cheaper analog equalizers and low-CPU digital plug-ins (which small-studio denizens usually have to rely on when dealing with high track counts) will frequently sound harsh and unmusical if too much boosting goes on. Much better to use cuts instead so that any nasty sounding distortion components are kept at a low level.

Now I'm not saying that you should never use EQ boosts, because there are some situations where it makes more sense to use boosts than cuts for mix-balancing purposes. For example, it's a lot easier to handle a peaking boost

LINEAR-PHASE EQ: PROS AND CONS

Most equalizers will inevitably change a sound's phase response as soon as you try to tweak the frequency balance, but that doesn't mean all of them do. An increasing number of digital "linear-phase" equalizers, for example, have been designed not to alter the phase of a processed signal at all, and these can provide more predictable control in situations where there is the potential for phase cancellation between mixer channels, such as when processing individual tracks in a multimiked recording. Linear-phase processing can also help when it proves necessary to apply heavy corrective EQ to a lead part, because normal "minimum-phase" EQ designs can make the processed sound less subjectively solid in this scenario, especially if a number of narrow EQ notches are active at once. However, it's important to realize that linear-phase EQ isn't some kind of cure-all—it's just as likely to introduce distortion and time-smearing side effects, so it needs to be used with the same degree of care. In particular, you should be on the lookout for any resonant filter ringing, because in linear-phase designs this usually precedes rather than follows each sonic event in the track—an effect that sounds particularly weird on drums and percussion. Linear-phase processing also tends to be much harder on your computer's CPU than normal EQ, so you may have to choose your moments carefully unless you want an ear bashing from Scotty in the engine room.

Although some DAW systems have linear-phase EQ available as a built-in plug-in, it isn't standard issue by any means. Fortunately, a number of third-party software developers provide linear-phase processing. DDMF's affordable cross-platform LP10 plug-in offers full control over the amount of phase shift each EQ band incurs, so you have the freedom to experiment with what sounds best in each case.

FIGURE 11.7
Linear phase equalizers, such as the IK Multimedia T-RackS device shown here, can be very useful when working with multimiked recordings, as long as you can handle the higher CPU munch.

than trying to implement a similar spectral change using a pair of shelving cuts. With some classic designs, the sonic side effects of boosting may actually be desirable and improve the subjective tone of the processed track. However, these situations are best treated as the exception rather than the rule. You need to exercise a great deal of caution wherever you use a boost, in order to minimize any degradation to the sound and to reduce the likelihood that you'll play mind games with yourself. In particular, try to keep peaking bandwidths

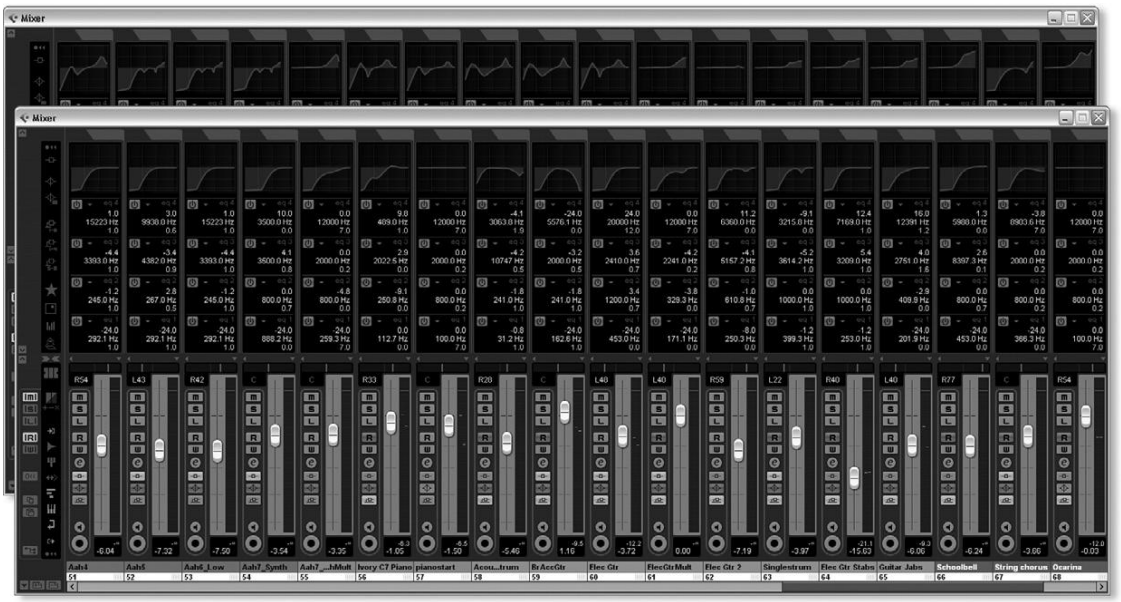

FIGURE 11.8
A before/after comparison of a range of EQ settings from one of my *Sound on Sound* "Mix Rescue" remixes. The addition of lots of EQ boosts on the original mix, especially at the high end, has added harshness, whereas the remix produced a smoother sound by concentrating on EQ cuts instead.

If you're boosting any track with more than one filter or with more than 6dB of gain, then you're probably overdoing it.

as wide as you can, and be wary of judging the effectiveness of any filter boost simply by switching it on and off. In addition I'd also suggest a general rule of thumb: if you're boosting any track with more than one filter or with more than 6dB of gain, then you're probably overdoing it. "If I'm having trouble with a mix," says Gus Dudgeon, "probably I'm using too much EQ. That's a pretty golden rule…. Basically, if you use too much EQ, you're trying to force something through a slot that's too small…. What [you] should be doing instead is backing off the EQ and turning the volume up."[14]

11.3 EQUALIZING MULTIMIKED RECORDINGS

Although the EQ techniques I've looked at thus far are fine for most simple recordings, there is an additional complication to contend with if you have multimiked (or DI-plus-mic) recordings among your multitracks: namely, if your EQ processing changes the phase relationships between the different mics and DIs in the recording, the outcome of the equalization may end up being as intuitive and predictable as a rollerblading platypus. "If you crank in a lot of EQ on the individual instruments, you really hear it," says Shawn Murphy, a specialist in mixing large-ensemble film scores, "and it alters the characteristics

of the leakage, so you're fighting the fact that in order to make the direct sound right on a certain mic, you've got badly EQ'd leakage to deal with."[15] So let's investigate this subject properly.

Multimiked Instruments

The simplest situation is where all the mics and DIs in the recording are capturing the same instrument without appreciable spill from any other instrument. With any luck you'll already have taken the opportunity at the balancing stage to buss all the component mic/DI tracks through a communal group channel to allow global control over them all, and you'll have set each track's relative level to achieve the best combination of timbre and balance you can. However, do double-check that this internal level relationship still holds water in the light of any processing you've done since constructing your first balance.

TONAL EQ

Although this chapter primarily centers around the balancing role of EQ, there is also a legitimate place in the mixdown process for equalization purely on subjective tonal grounds—in other words, equalizing a sound to make it more attractive, rather than just to fit it into the mix balance. However, there's not a tremendous amount of useful advice anyone can give for this kind of EQ work. Just crank knobs until it sounds better!

It's in this context that classic analog equalizer designs really come into their own, because the processing by-products of these are often as much a part of their appeal as their ability to sculpt frequency response. Which equalizers suit which tracks is a subject that frequently seems to lead to pistols at dawn among different practitioners at all levels, so only your own experience can really guide you. Experimentation is the key to success here, and there's little excuse not to have a go now that digital emulations of classic analog devices are so numerous. My advice is to make a habit of trying out two different "character EQ" plug-ins every time you're looking for extra tonal magic. Do your best with each one, and then switch between them to assess their relative merits before choosing your favorite. This doesn't take too much extra time to do, but it quickly builds up your own knowledge of how different classic EQ designs respond to different sound sources.

Oh, and feel free to completely disregard my tutting about boosts when you're using EQ solely for tonal coloring (rather than for balancing), because they'll alter the attitude of a sound more than cuts will, especially when you're using emulations of classic EQ units. You might also want to chain several contrasting EQ plug-ins together on a single track, and that's all right too as long as you're convinced that the sound (and not just the loudness) is improving. "I think you can get [better results] by not trying to get everything out of one EQ," says Dave Pensado. "Every EQ is good for certain things. Even the cheapest EQ has a use. Don't judge gear by its cost or inherent quality; judge it by its uniqueness."[16] However, when chaining EQs do try to be clear in your mind whether each band of processing you use is for tone or balance, because you're likely to come a cropper if you try to balance your mix with massive character-EQ boosts or attempt to overhaul a sound's subjective timbre using bland balancing-EQ cuts.

If the group channel fader of the instrument still feels unstable in a way that appears to demand EQ, then you've got three options available:

- You could tamper creatively with the timing, phase, and polarity of the individual mics to achieve a more suitable frequency balance in the combined sound. The upside to this option is that you can achieve dramatic changes if necessary without incurring any processing artifacts, because no actual processing is involved. The downside, though, is that this process is very hit and miss—you can experiment all you like in the hope of finding a better sound, but there's no guarantee you'll find it.
- You could EQ the individual mics/DIs in the setup separately to achieve the best balance of desirable frequency components from each. If you have a target balance in mind for the instrument, it may help to listen to each of its component tracks separately to ascertain which track is most appealing for each frequency region, so that you can use EQ to favor the best bits on each track. However, this approach may not really bear fruit unless you've minimized any comb filtering between the tracks (so that they seem to add together fairly predictably) and the EQ processing doesn't unduly alter the phase response. Although you sometimes get lucky, in practice I've found that this kind of equalization inevitably ends up being a bit of a lottery most of the time and often necessitates remedial phase/polarity adjustments.
- You could EQ the group channel that carries the combined instrument signal. This will be the most predictable of the three options, because you won't skew the phase relationships between the individual mics/DIs—you can process it just like a single mono/stereo recording, as described in Section 11.2. On the other hand, though, this method also gives you much less scope to alter the instrument's balance or subjective timbre.

Clearly, there are many ways to skin this particular cat, and it's fair to say that each individual engineer will have preferred working methods here, which will also change depending on the exact nature of the recording. The only real guidance I can give therefore is to say that a blend of all three tactics tends to bring about decent results most efficiently. So I might start off tinkering with timing, phase, and polarity adjustments to get the overall character and tone of the instrument into the right kind of ballpark, but without spending all year on it; then I might apply EQ to massage the balance of each mic/DI within the most important frequency ranges, favoring settings which don't impact too heavily on the phase response; and finally I might EQ the instrument's global group

EQ FOR DISTANCE

Because high frequencies are so easily absorbed in nature, this affects the frequency response of sounds arriving from a distance. Tony Maserati explains how he learned to use this psychological effect to advantage at the mix: "I learned from listening to Roger Nichols and Steve Hodge…. They would use EQ to move things front and back. Obviously, the brighter something is, the closer it is; the duller it is, the farther away."[17]

channel to deal with masking of, or by, its frequency response within the mix as a whole.

Multimiked Ensembles

Although there are no additional EQ tactics for multimiked ensembles beyond what we've just looked at in relation to multimiked instruments, there are further decisions to be made about which tracks are to be EQ'd separately and which are better processed in combination. In a drum kit recording, for example, you might wish to group three snare drum mics together to gain global EQ control over their combined frequency balance, while using separate phase/polarity adjustments and equalization on each individual mic to determine the exact nature of that timbre; but you might then further EQ the snare sound from the main drums group channel where your adjustments won't affect the snare microphones' phase relationship against the overheads. At the end of the day, your ears must decide on the most successful outcome, and even the professionals disagree as to the best approach. For example, many rock engineers think nothing of processing individual instrument mics in an ensemble, but veteran producer Renaud Letang expresses concerns with this approach: "I find that if you apply EQ or compression to individual [instrument] parts you lose the impression of… a unified section, and it becomes really easy to lose the overall balance. But when you pre-mix and then EQ or compress the whole section, the sound remains much more natural."[18] My advice is just to keep a cool head, add in the mics/DIs as you did during the first balance, and listen carefully for unwanted knock-on effects whenever you process by muting and unmuting the track or group in question.

There's a limit to what EQ can do against some tonal and balance problems, for two reasons: it doesn't actually add anything to the sound, and it can't deal with frequency imbalances that alter over time.

One other point to make is that you may wish to reassess the high-pass filter settings you initially decided on while balancing, in the light of your increased knowledge of the mix. If you remember, we set them fairly conservatively to start with in order to defer dealing with additional phase problems. Now that we're better equipped to handle those issues, you should feel free to clear out unwanted low end more astringently should you so desire.

11.4 THE LIMITS OF EQ

For mixing purposes, that's pretty much all you need to know about EQ. However, there's inevitably a limit to what it can do against some tonal and balance problems, for two reasons. First, it doesn't actually add anything to the sound—it just adjusts the levels of the frequencies already present. For example, a dull-sounding electric piano feeding a beat up keyboard amp is never going to sound very bright, no matter how much you boost the high end, because there's nothing much happening above 4kHz other than noise.

The second problem with normal EQ is that it's static, so it can't deal with frequency imbalances that alter over time.

Again, don't spend weeks trying to solve every mix problem with EQ. Be happy to improve things as far as you can, and leave any trickier troubleshooting to more advanced processing. My own motto is this: if I've spent more than five minutes equalizing a track, then EQ is probably not the whole solution to its problems.

CUT TO THE CHASE

- The main purpose of equalization at mixdown is to address frequency-domain problems with the mix balance. Although there's a place for more subjective tonal EQ'ing at mixdown, this should be a secondary consideration.
- Add your instruments into the balance in order of importance, so that it's more straightforward to deal with frequency masking between different tracks. Always confirm the validity of your EQ judgments within the context of the mix, rather than in solo mode. Avoid graphic EQ designs for mixing, and reach for shelving EQ before peaking EQ. Favor cuts over boosts wherever possible. If boosting is unavoidable, then do your best to limit it to one EQ band, keep the gain below 6dB, and use as wide a bandwidth as you can get away with. Try to remain sensitive to both wide-bandwidth and narrow-bandwidth frequency-balance issues. To maintain your objectivity, bypass each filter frequently as you set it up, and make use of different monitoring systems and listening levels. Once you have an ideal balance for a processed track in its own right, switch it in and out of the mix while listening for unacceptable masking of other more important sounds, applying further EQ cuts at conflicting frequencies if necessary.
- When equalizing multimiked recordings, you need to bear in mind that EQ may alter the tone and balance of the sound unpredictably because of changes in the phase relationship between the individual mic/DI tracks. Applying EQ to individual tracks will give a different outcome than processing group channels. Linear-phase EQ may help increase your options here, but it is not without its own processing side effects.
- If a track's dynamics plug-ins aren't doing the job properly, then equalization earlier in the processing chain may improve the musicality of the results. Postdynamics EQ, on the other hand, is a better bet if you're already happy with the dynamics. It's not uncommon to need both pre- and post-dynamics EQ.
- When the main purpose of an equalizer is to make a sound subjectively "better," rather than just to balance its frequencies in the mix, the rule book pretty much goes out the window—no setting is wrong if it achieves a sound you like. Make a point, though, of trying out two different equalizers every time you're equalizing with this intent, so you begin to build experience concerning which models suit which types of sound. Also be sure you're not fooling yourself into thinking something sounds better just because boosts are making it louder.

Assignment

- Track down all the equalizers on your DAW system so that you know what options are available. If you have little choice, then consider supplementing your selection with third-party plug-ins. Try to get hold of a linear-phase equalizer too if you work with a lot of multimiked recordings.
- Mute all the tracks in your mix and rebuild the balance again, this time applying equalization in search of a more stable balance.
- Make a note of any faders that still feel unstable.

www.cambridge-mt.com/ms-ch11.htm

This page intentionally left blank

CHAPTER 12

Beyond EQ

As I said in Chapter 11, equalization is extremely useful for mixing, but it can only adjust the balance of frequencies that are already in the recording. Much of the overprocessing in typical small-studio mixdowns arises from people maxing out their EQ settings in a misguided attempt to add something new to the sound. So if EQ isn't giving you a sound you like or is unable to deliver a satisfactory mix balance, then it's time to stop footling about with what's already there and actually add information to the mix instead.

12.1 DISTORTION AS A MIX TOOL

The textbook situation in which EQ has nothing to offer is where you want to brighten a bass guitar/synth line with no appreciable high frequencies. This might be because the synth's low-pass filter has cut away all the high end or simply because a limited-bandwidth oscillator waveform was used to start with. In such a scenario, there may well be no energy at the upper end of the frequency spectrum, so equalization is powerless to alter the subjective tone. The only alternative is to create new high-frequency components that weren't originally there.

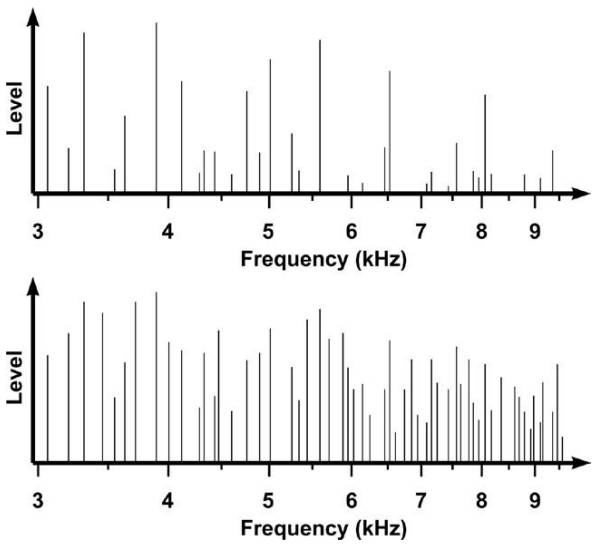

FIGURE 12.1
The upper graph has less harmonic density than the lower graph. If the harmonics of an instrument aren't dense enough, then distortion may help thicken it up.

That's just the most obvious case where EQ doesn't deliver, though, because even where an instrument does generate frequency components in a required spectral region, they may not be dense enough to give that spectral region a sense of fullness within the balance. In this kind of case, you'll be unable to set a satisfactory EQ boost. It'll either feel too small, because the frequency region isn't dense enough to compete with other sounds in the mix, or it'll feel too big, because the few frequencies that are in that spectral range are poking too far out of the balance. Increase the

harmonic density, however, and that spectral region can remain solid in the mix without individual frequency components sticking their heads over the parapet.

The tool that can help in both these scenarios is distortion, an effect widely used in professional mixing circles but pretty much ignored in a lot of small studios. Although people take great pains to minimize distortion a lot of the time in record production, that's no reason to neglect its creative possibilities at the mix. The bottom line with distortion devices is that they add new harmonics to the processed signal, and because the frequencies of these are usually related in some way to the frequencies of the unprocessed signal, they retain a musical connection. In short, distortion allows you to change the timbre of a track without substantially altering its musical function in the arrangement. As such, its use is almost as widespread as EQ, and Rich Costey is not the only one who regularly distorts even lead vocal tracks: "[It's] so commonplace nowadays that to me it's the same as distorting a guitar or a bass. And much of the time when I'm mixing records for other people I'm distorting the vocal, whether they know it or not!"[1] Probably the most commonly name-checked device for mix distortion is Tech 21's SansAmp range, which can find its way onto just about anything in modern mix work, especially if your name's Tchad Blake: "I actually have one for bass, one for kick, and one for snare…. It's funny—a little distortion goes a long way."[2] The big challenge when using distortion in a mix setting, though, is restraining it so that it adds harmonics only at those frequencies you want it to. Here are some tips for keeping it under control.

First, it's nice to have different distortion devices to choose from, so that you can mix and match. There are lots of electrical devices that can create musically useful distortion, and many of these have now been digitally modeled for use in plug-in form: tape machines, guitar amps, effects pedals, soft-clippers, vacuum tubes, transistors, and transformers, to give but the briefest of lists. Many of these plug-ins also have selections of presets that offer widely varying sounds, and you don't need to understand anything about how any of them work just to browse through and find one that seems to add something useful to the track you're processing. Even a lo-fi digital "bit-crusher" distortion can occasionally be useful. "What that does is create some contrast between that track and the rest of the song," says Serban Ghenea. "You don't need much of it, but it can make a world of difference."[3] Be aware, though, that the digital distortion products from these devices tend to be less closely related to

FIGURE 12.2
A selection of different freeware distortion flavors (*top to bottom*): rectifier distortion (GVST's GRect), variable-knee clipping (GVST's GClip), tape saturation (Jeroen Breebaart's Ferox), valve distortion (Silverspike's Ruby Tube), and transformer distortion (Bootsy's Tessla SE).

the input signal on a musical level, so be on the lookout for uncomfortably discordant added harmonics, especially on pitched instruments.

Second, some distortion processors have loads of controls, but to be honest life's too short to spend ages learning what most of them do, so you might as well just laugh maniacally and twiddle with abandon. What actually matters more than anything with distortion devices is how hard you drive them. If you hit a distortion process gently, you may just get a subtle thickening of the sound, whereas if you slam your track in at high level, you may barely be able to recognize the original sound amongst all the added frequency-domain hash. If there's any control on your distortion unit labeled Drive or Gain, then that's the one you should pay most attention to. Occasionally no such control is provided, in which case you can vary the degree of distortion by adjusting the signal level before the distortion device's input.

MORE ADDED HARMONICS: HIGH-SPEED MODULATION EFFECTS

Another way you can add harmonic density to a signal is to treat it with audio-frequency modulation effects. There are a number of different possibilities to choose from, including frequency modulation, amplitude modulation, ring modulation, and any vibrato/tremolo/auto-wah plug-in that can set its modulation frequency above 20Hz. If you fancy getting to know the mechanics of how these effects work, go ahead, but there's no obligation—you can use them just fine in a mix situation by surfing the presets. The only really important control you need to know about will be called something like Modulation Depth, and it basically determines the amount of frequency information added.

What you'll notice pretty quickly, though, is that frequency additions from these effects aren't musically related to the input signal, so you'll usually introduce dissonant tones into pitched parts. Although simple melodic parts can nonetheless still be coaxed into responding musically if you experiment with the effect's modulation frequency, in practice you're more likely to get useful results with unpitched sounds. That said, the very unmusicality of some of these effects can be a boon on occasion. For example, I've used ring modulation to add growl to aggressive lead vocals on more than one occasion.

I'm frequently amazed by how effectively you can remodel or enhance a lackluster track simply by browsing a few different distortion plug-ins, choosing your favorite, and then adjusting the drive to taste. It's so simple that it almost feels like cheating! However, there are significant limitations with this basic insert-effect configuration. First, increasing the distortion's drive will simultaneously increase both the level and the density of the added harmonics, whether you want that or not. Second, you can't directly specify how much

the distortion products in each spectral range balance with the track's original frequencies. For these reasons, a parallel processing approach (in other words, using the distortion as a send effect) is usually my preferred choice, because this provides much more scope for deciding exactly what it is the distortion adds. You might, for example, decide to drive the distortion hard to create denser harmonics, but mix those in with the unprocessed track at only a moderate level. You are also free to EQ the distortion channel to funnel the added spectral energy into those areas where it's really needed. You might also consider compressing the signal before it reaches the distortion to maintain a more consistent density of added harmonics. As with all parallel processing, though, you do have to be careful of phase mismatches between the processed and unprocessed channels (particularly when using guitar amp simulators). I often end up using phase rotation in such cases to get the best combined sound, especially if I've resorted to EQ to carve away unwanted regions of the distortion sound.

High-frequency enhancement can be tremendously fatiguing on the ear, and your hearing will quickly become desensitized to it.

One type of parallel distortion processing warrants special attention, and that's the high-frequency enhancement technique pioneered by Aphex in its Aural Exciter. In simple terms, with this type of parallel distortion technique the signal feeding the distortion is aggressively high-pass filtered and then compressed before it reaches the distortion. Mixing this processed signal in with the unprocessed track gives an unnatural level of high-end density, which can help sounds seem detailed, airy, and more upfront in the mix.

Needless to say, the Aphex process has long been a common choice for lead vocals. If you're going to use this effect, however, there are a few issues to beware of. Most of all, you need to realize that this effect can be tremendously fatiguing on the ear, and your hearing will quickly become desensitized to it, so it's more important than ever to keep checking your objectivity. Switch the effect out of circuit, allow your ears to adjust to the preprocessing balance, imagine the sound of the processing you're after, and then pop the enhancement back into the mix to check that it's actually doing what you want it to. If you don't take this precaution, you'll fall into one of the oldest traps in the book: turning up the enhancement too far because your fatigued ears have become less sensitive to it and then discovering after a night's sleep that your singer actually sounds like a wasp shaving!

Another point to bear in mind with the majority of distortions is that they can reduce dynamic range, so you may need to apply additional balance processing to compensate. In many situations, though, this side effect can actually be beneficial, particularly when processing noisy signals. For example, I frequently find myself using clipping to round off over-spiky percussion transients.

PITCH-SHIFTING FOR DENSITY

Another way you can add harmonics to an instrument is to use a pitch shifter to duplicate it into a different pitch register. With melodic and harmonic parts, the most foolproof interval will usually be an octave, but major/minor thirds (shifts of three/four semitones) and perfect fourths and fifths (shifts of five and seven semitones) can also work well in some instances. Upward shifts tend to be more useful than downward ones, and you'll typically want to mix the shifted channel at a low enough level that it's perceived as part of the unprocessed track, rather than as a new set of notes in its own right. This technique is usually less successful with drum and percussive sounds, because signals full of transients and noise tend to stress pitch shifters beyond the point where they sound any good. That's not to say that the idea's a dead loss, but the results are likely to be patchy.

12.2 LOW-END ENHANCEMENTS

By its very nature, distortion processing tends to add most, if not all, of its added energy above the lowest frequency of the unprocessed sound, and when you're looking for tonal changes, this is usually perfectly appropriate. However, there are some situations where you might want to add "oomph" at the low end of an instrument's frequency range, or even below that. This is where a type of processor called a subharmonic synthesizer can help by artificially generating additional low-frequency energy that is still musically related to the unprocessed sound.

FIGURE 12.3
A couple of subharmonic enhancer plug-ins: Waves LoAir (*top*) and Apple Logic's bundled SubBass (*bottom*).

Although high-profile engineers such as Jack Douglas,[4] Jon Gass,[5] Jacquire King,[6] Alan Moulder,[7] and Thom Panunzio[8] have all professed themselves users of such processors, I have to be honest and say that it's rare to hear anyone applying them effectively in the small studio. This is partly because you need good low-end monitoring to gauge their true impact, but it also has to do with the fact that every manufacturer of a subharmonic synthesizer appears to use a slightly different algorithm, and you can never be quite sure how well each will respond to a given bass-instrument recording. One plug-in might work great on kick drums but churn out a load of woofy rubbish when presented with a bass line; another might handle only midregister bass parts; another might handle a good range of bass notes but throw a tantrum if there's any drum spill on the bass mic.

Therefore, I personally rely on alternative strategies most of the time. For kick drums I insert an additional sample, whereas I supplement melodic bass instruments with a MIDI "subsynth." Not only do both these techniques allow

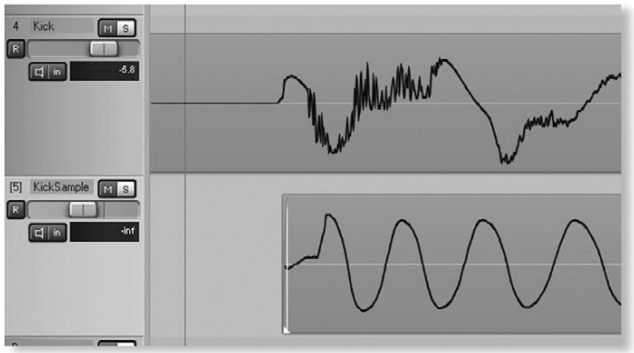

FIGURE 12.4
If you're layering a sample against a live kick drum, you need to make sure that the polarity of the sample matches the live track for every hit. In this screen you can see an instance where it's not matched, and this results in a noticeable loss of low-end weight.

complete pitching and timing reliability, but they also offer much wider scope for tailoring the timbre of the enhancement to suit the production as a whole. Let's deal with each in turn.

Drum Triggering

There are now lots of software plug-ins that will automatically detect drum hits in a track and play back a sample in time with them, but with kick drums you'll only get good results if you refine the placement of each triggered sample by hand. "Software is very time-saving, but the drawback remains that it's not as good as doing it manually," says Cenzo Townshend. "You still have to go through each hit and move it in time to make sure it's 100 percent correct and the phase is correct. You look at the waveforms, but you also need to use your ears."[9] Even once you've made sure that the phase relationship between the original and triggered kick tracks remains fairly constant, though, it's as well to investigate whether additional overall timing/phase adjustments will further improve the combined sound.

Of course, sample triggering can also be used much more widely than this, allowing you to add new characteristics to any drum sound. It's particularly useful with live drum recordings which, for whatever reason, need to rely heavily on the close-mic signals—for example, if the overhead mics are picking up mostly cymbals. Close mics positioned just a few inches from each drum will spotlight only a very small portion of the instrument's complete sound, so the recordings from these mics will usually be rather unnatural and uninspiring to listen to. (Certainly, nine out of ten snare close-mic recordings I hear from small studios just go "donk"!) Triggering a sample alongside, or instead of, the close mic can help to provide a more pleasing overall tone in such cases. Andy Wallace uses an interesting variant on this technique to supplement his snare ambience: "I use the samples to drive reverbs…. The thing I like is that I can EQ them so that I can really tune the ambience and where it sits in the whole frequency response…. More so than I can with the overheads, because I normally EQ those so that the cymbals sound the way I want them to sound…. [I can] shade with a little more control using my ambient sample."[10]

But there's another good reason for triggering samples within the context of a live drum kit recording: triggered samples have no phase relationship with the other microphones and no spill. This means that you can freely balance and process the sample without fear of destroying the rest of the kit sound—great if you only need the sample to supplement the sound in a specific way. For example you might apply slow-attack gating and assertive high-pass filtering to isolate just a snare sample's high-frequency sustain components.

FIGURE 12.5
If you don't have suitable drum samples for mix triggering, then it's easy to lay hands on commercial sample libraries from outlets such as Time + Space.

The Internet being what it is, it's child's play to find a selection of decent samples for most drums. However, on the offchance that typing "free drum samples" into Google doesn't provide more choice than you could ever want, there are lots of commercial outlets that can provide you with tons of material. Check out *www.timespace.com* or *www.soundsonline.com*, to mention but two large-scale international distributors.

Incorporating a MIDI Subsynth

Propping up a lightweight bass line with a low-frequency synth part is a stalwart of many urban styles in particular. "The bass is very important," explains DJ Premier. "If the low end of a [loop] sample isn't really heavy, I'll always follow the exact bass line of the song and put that underneath. A lot of people ask me what EQ I use to get the bottom end of my samples to come through so strongly, but I'm like 'Man, it's not EQ! I'm playing the same notes verbatim.'"[11] This is often pretty straightforward to do if the bass is repetitive and tightly locked with a regular song tempo, but if not then there may be some initial donkeywork involved in sorting this out. Even if it takes half an hour, though, that time investment is usually worth the peace of mind that comes

from knowing that the low-end frequency balance will be totally dependable. Usually only the freakiest of extended jazz-funk free-for-alls have me waving the white flag and scurrying back to traditional subharmonic processing.

Once the MIDI part is in place, you need to consider what kind of sound the subsynth should use. One common option is to double the bass line at the octave below using a simple sine-wave patch. In this case, there's nothing much to do other than set the sine wave's level and listen to it in solo just to check that the synth's envelope settings aren't so fast that they create unwanted clicks. Things get more complicated if you're using a subsynth to try to beef up the existing fundamental frequency of a bass part, because the subsynth will overlap the main bass part's frequency range. If the subsynth's waveform isn't in phase, then phase cancellation might actually cause the combined sound to become even skinnier! Moreover, because this phase relationship will usually change from note to note, you could end up with an uneven low end that's all but impossible to balance with the rest of your mix. If the existing bass sound's fundamental frequency is strong enough to cause difficulties such as this, then try to get rid of it to make space for the subsynth at the low end. A steep high-pass filter on the main bass part is one solution here, but at times you may need to use the more surgical approach of notching out individual frequencies, preferably with the assistance of a high-resolution spectrum analyzer.

In the event that you want to do more than just add a suboctave or emphasize the existing bass fundamental, a sine-wave synth oscillator may not be the best bet; other waveforms will fill out the overall tone better. I like using a triangle wave most of the time, because it doesn't have such dense harmonic spacing as a sawtooth and it's duller-sounding and less characterful than a square wave—both of these properties seem to make it better at blending with (rather than overwhelming) the sound it's layered with. Whatever waveform you use, though, you should also take the same precautions against phase cancellation of the bass part's fundamental frequency. I'd also steer clear of detuned multioscillator patches, because the "beating" between the two detuned layers may cause the subsynth's fundamental frequency to fluctuate unacceptably in level. It probably makes sense to stick with mono patches as well so you don't get mono-compatibility problems. These restrictions mean that you only really need a simple instrument to generate subsynth parts. Whether you engage the subsynth's own filter or sculpt its output with EQ is purely a question of what kind of low-end tonal enhancement you're looking for. With a triangle wave in particular, you might not feel the need to filter it at all, although I do personally employ some kind of gentle low-pass filter most of the time to restrict its contribution to the lower octaves.

FIGURE 12.6
When designing subsynth parts, your choice of synth waveform is an important decision.

One final point to make is that subsynth parts usually need to be controlled tightly in terms of dynamic range

COMB FILTERING AND RESONANCE AS TONAL TOOLS

Another method of adjusting the subjective tone of a track is to deliberately introduce the complex frequency-response undulations of comb filtering by mixing in a delayed version of the signal. All you do is set up a delay plug-in as a send effect and feed it from the track that needs timbral adjustment. Choose any old preset you like, and then set the effect's Delay Time parameter to anything less than about 20ms. This should result in obvious comb filtering when the delayed signal is mixed in with the unprocessed track. The way the filtering sounds will depend on the exact preset and delay time you've used, so tweak both to taste while keeping a hand on the delay return channel's fader to control the severity of the filtering. Some of the presets may introduce pitched sympathetic resonances, and if these are sensitively tuned using subtle adjustments of the Delay Time control and then subtly applied, they're capable of adding extra musical sustain. Always keep in mind that you can EQ the delay return too, should you wish to curtail the comb filtering and sustain artifacts in some areas of the spectrum.

Related to this kind of short delay effect are the three common modulation treatments: chorusing, flanging, and phasing. Fun as they are for special effects, from a mix perspective they don't really serve much of a technical purpose, so you might as well just surf the presets and tweak things erratically in search of something you like. (Some people say that chorusing can create the illusion of a larger ensemble, but it never seems to work like that for me!) They can be useful for adding stereo width in some instances, though. (See Chapter 18 for more on that subject.)

Artificial reverb send effects can also be coaxed into producing a wide range of changes to tone and sustain. Again, feel free to process the reverb return channel, and consider in particular narrowing its stereo image if you're processing mono sounds, so that the added width doesn't reveal the effect's synthetic nature. No need to fuss with understanding the controls at this point, just look for a preset that sounds good. Each will have a unique tonality, echo pattern, and overall resonant character, so you may be surprised how different they all sound in this context. Don't ignore any rubbishy reverb plug-ins either, because they often yield the most dramatic timbral changes. As with many mix effects, reverb doesn't have to sound any good on its own as long as it combines musically with the track you're processing. (If you want to get deeper into using tonal delay/reverb effects, skip ahead and browse Chapters 16 and 17.)

or else they can really eat into your track's overall headroom. This is where the MIDI-triggered method really pays dividends, because it's simple both to restrict the dynamics of the MIDI part and to compress the subsynth's output. So even if your low-frequency monitoring facilities aren't up to much, you can still rely on a subsynth to give consistent low end if you pay attention to your DAW's meters.

12.3 SYNTH PADS

Use of MIDI sound sources for mixing purposes doesn't stop at subsynth parts, by any means, because chordal synth pads in particular can be immensely

The first mistake most newcomers make is neglecting to equalize a synth pad's frequency response, because it's vital that pads only operate precisely where they're required.

useful across a surprisingly wide range of different music styles. Despite their inglorious history, pads don't have to turn your mix into some kind of 1980s bubble-perm nightmare as long as you design and apply them effectively. Indeed, pads can easily add warmth and sustain to thin-sounding acoustic guitar parts, for example, or improve the pitch definition of mushily overdriven electric guitars, without anyone being the wiser. There are very few parts that you can't supplement with a pad synth if needs be, just as long as you pare each pad down to its bare essentials—that's the key to keeping pads from becoming too audible in their own right or trespassing unnecessarily on other aspects of the production sonics.

The first mistake most newcomers make is neglecting to equalize the synth's frequency response, because it's vital that pads only operate precisely where they're required. The exact note pitches and synth timbres also warrant careful thought, as these can dramatically affect the density of harmonics within the target frequency ranges. Pads don't have to be long held chords, either, because in many cases a simple rhythmic repetition of synth chords is better at reinforcing the track's groove. When it comes to deciding on a fader level for the pad, you don't often want the general public to notice you're using a pad at all—bands who specialize in natural-sounding acoustic styles or hard-hitting rock music don't really like to brag about their synth chops, for instance, despite the improvements that pads can bring to guitar textures in both those genres. So my normal rule of thumb is this: if you can hear the pad, it's too loud! To put it another way, you should only notice that the production has a pad in it when you toggle that channel's mute button.

CUT TO THE CHASE

- Equalization can only adjust frequencies that are already there; if you need to add something new to the sound instead, then you'll need different processing.
- Distortion is something of a secret weapon for filling out small-studio productions. You can achieve quick mix improvements using distortion effects as insert processors, but it's usually preferable to use a parallel processing configuration instead. Just be wary of unwanted comb filtering should your distortion plug-in tamper with the processed signal's phase response.
- Pitch shifting can add complexity to a track's frequency response if you take care with the shifting interval used. This tactic usually sounds best for simple pitched instruments and upward shifts. High-speed modulation effects are another option, although care must be taken with pitched instruments if you wish to avoid unmusically dissonant frequencies being added. Short delays and reverbs can deliver tonal changes through complex comb

filtering and have the potential to add subtly appealing sympathetic resonances in addition.

- Subharmonic synthesizer plug-ins are useful on occasion, but adding supplementary sample/synth parts is usually a more successful way to add low-end weight to bass instruments.
- Added drum samples not only provide the opportunity for timbral overhaul, but they also allow independent manipulation of one instrument within a multimic recording without compromising spill from other sources. Although automatic drum-triggering software is available, you should refine the timing of each sample by hand to avoid phase cancellation against the trigger track.
- MIDI-driven synthesizers are excellent tools for adding spectral complexity. Subsynth patches can help provide small-studio productions with powerful yet controllable low end, whereas pads can usefully fill out the sound of almost any pitched instrument, even in musical styles where you'd least expect it. It's rarely desirable to hear a pad in its own right, so you should use your MIDI programming, sound-design, and equalization skills to tailor its frequency additions to the specific task at hand.

Assignment

- Do a quick survey of the processing available on your mixing system to get an idea of your options with respect to the various different effects mentioned in this chapter: distortion, high-speed modulation, pitch-shifting, subharmonic synthesis, delay, and reverb.
- Investigate your DAW's facilities for triggering drum samples in response to recorded audio, and for generating MIDI parts to shadow specific pitched instruments.
- Start building a library of good drum samples for mixing purposes so that you don't have to hunt around for ages every time you need one. Also check out what synths you have available to you, and make a note of any presets that turn out to be more useful than others in practice.
- Mute all the tracks in your mix and rebuild the balance, addressing any frequency-balance concerns that careful equalization was unable to resolve.
- Make a note of any faders that still feel unstable.

www.cambridge-mt.com/ms-ch12.htm

This page intentionally left blank

CHAPTER 13

Frequency-Selective Dynamics

Chapters 9 and 10 demonstrated various ways to handle time-domain balance problems, and Chapters 11 and 12 focused on frequency-domain issues. Although those techniques alone may be enough to deliver you a workable balance, usually at least one instrument will remain that can't be persuaded to maintain its proper place in the mix by these means. This is because some balance problems cannot be overcome unless you process both the time and frequency domains at once. To put it another way, not one of the processing methods we've discussed so far is much use against a dynamic-range problem that exists only at high frequencies—a normal dynamics processor's gain reduction won't target that frequency region separately, and the static frequency-balancing facilities of equalization won't tackle the level inconsistencies.

To pull off this kind of stunt, you need frequency-selective dynamics processing. Plug-ins that provide this feature do tend to look a bit intimidating, so it's hardly surprising that many small-studio operators leave them well alone, whereas others just pick a preset and hope for the best. However, these devices are too powerful to ignore if you want commercial mix results, but they are also powerful enough to completely wreck your tracks if applied in a scattergun fashion. This chapter shows how you can combine time-domain and frequency-domain treatments for yourself to tackle some of the toughest balancing tasks that small-studio productions face.

Frequency-selective processors are too powerful to ignore if you want commercial mix results, but they are also powerful enough to completely wreck your tracks if applied in a scattergun fashion.

13.1 FREQUENCY-DOMAIN TWEAKS FOR FULL-BAND DYNAMICS PROCESSORS

Before I introduce any new types of processor, let's have a look at how the dynamics devices we've already covered (often called "full-band" processors,

Mixing Secrets for the Small Studio.

because they operate on the entire audio spectrum) can be coupled with EQ to increase your balancing options.

Equalizing Parallel Processors

One simple way to begin exploring how dynamics can be aimed at specific frequency regions is by elaborating on an idea we've already dealt with quite a bit: parallel processing. If you set up any dynamics processor as a send effect, then you can EQ its return channel to tailor the processed tone. So where a fast-attack/fast-release compressor working in a parallel configuration would normally just emphasize overall sustain, a low shelving cut to the return signal would direct that sustain more toward the upper frequency regions—a good way of compensating for the shorter sustain of higher piano notes, to give one example. Dave Pensado supplies another application, this time for the kick drum. "You can have the main kick drum fader, and another fader where you compress the dogsnot out of the kick, and add a lot of top end to it. Then, maybe in the chorus you add in a little bit of that second chain. It gives a tiny bit more attack… as if the drummer was hitting it harder."[1]

Dylan Dresdow combined compression, EQ, and gating to achieve frequency-selective transient enhancement for his work with Black Eyed Peas: "I'd mult the kick drum and mult the snare, compress them really hard, and roll off a lot of high end and low end, so they become very midrangey, and I'd put a tight gate on them. I'd sculpt the attack in this way and would add this into

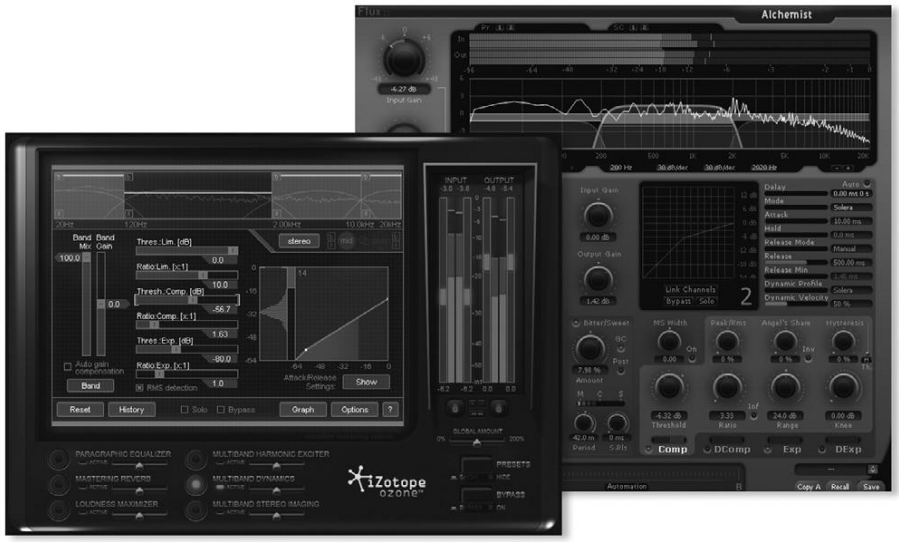

FIGURE 13.1
Complex frequency-selective dynamics processors such as iZotope Ozone or Flux Alchemist may appear intimidating, but there's actually no great mystery to using them if you remain focused on the central issue of mix balance.

the track."[2] Of course, you might achieve a similar kind of effect by EQ'ing a parallel transient enhancer. In all of these situations, a little further EQ may be desirable on the combined sound to compensate for any overall level increase in the enhanced frequency region—but that shouldn't be a big deal, because you already know exactly where that region lies.

As usual in parallel setups, you do need to be careful that EQ phase shifts don't cause any unpleasant comb-filtering side effects when the processed and unprocessed signals are combined—this happens often enough for me that I now use linear-phase EQ for parallel dynamics return channels as a matter of course. There's also the question of whether to use the equalizer before or after the dynamics, although this should hold no new mysteries if you've already absorbed the advice in Chapter 12: if the gain changes seem musical, then EQ after the dynamics; if not, then try using EQ earlier in the processing chain. Returning to my earlier piano example, if you want to add extra high-frequency sustain with a filtered parallel compressor, it makes sense to put the filter before the dynamics. That way the compressor stamps on the high-frequency signal peaks from high-pitched notes and smoothly resets its gain reduction between them to allow the interpeak sustain tails to shine through. Had the high-pass filtering been placed after the dynamics, the lower notes would also trigger gain reduction, causing unnatural lumps and bumps in the high-note sustains.

Refining Your Gating Action

Another common hybrid technique is where equalization is applied to the dynamics processor's level-detection circuit (or "side chain"), but not its main signal path. This allows you to adjust the sensitivity of the processing to different frequencies without directly changing the tonal balance of the processed audio. You're most likely to find side-chain EQ facilities on gates (usually in the form of simple high-pass and low-pass filters) where they can help a great deal in getting the gate to open and close the way you want it to. Without such facilities it can be well nigh impossible to reduce unwanted spill signals in multimiked recordings of closely spaced instrument ensembles such as drum kits.

FIGURE 13.2
Like a lot of well-specified gates, Steinberg Cubase's bundled Gate plug-in features built-in side-chain equalization to refine its triggering action.

Take the example of a microphone positioned underneath the snare in a drum kit. This will normally pick up the kick drum along with the snare, so if you only like its snare contribution, then it makes sense to gate out everything else. The problem is that the level of spill from the loudest kick drum hits may exceed the level of direct sound from the softer snare hits, so no setting of the gate's level threshold will get it working how it should: either some of the kick spill will make it past the gate, of you'll lose a few of the softer snare hits. Side-chain EQ can improve this situation by filtering the low end out of the signal

feeding the gate's level detector, thereby desensitizing it to the powerful low frequencies of the kick drum spill. The great thing about the EQ being only in the side chain is that you can slap it on with a trowel without affecting the tone or phase response of the processed snare sound at all.

Other applications of side-chain EQ for drum kit gating abound: reducing hi-hat spill on the main snare signal, desensitizing the gate on a floor tom to its own rumbling sympathetic resonances and to accents on nearby cymbals, and drying up the room ambience on an out-the-front kick mic. However, there is one thing to be careful of: if a drum's attack transient is mostly made up of high frequencies, then filtering those out of the detector side chain may cause the gate to open slightly late, making your kick all icky.

Sibilance Reduction

Side-chain equalization of compressors is slightly less common but can nonetheless be a godsend when you're trying to minimize the unwanted side effects of heavy compression. This is nowhere more relevant than with vocals, which have a wide inherent dynamic range that's frequently pulverized in many mainstream music styles for the sake of maximum audibility or attitude. Because most compressors are, by nature, less sensitive to high frequencies, the most common unwanted side effect of heavy vocal compression is that sonic events that contain mainly high frequencies tend to cause a lot less gain reduction and therefore end up being overemphasized by the processing. With vocals, that means that breath noises and consonants such as "ch," "f," "s," "sh," and "t" will all tend to rise up too far in the balance—particularly the "s" sounds (or sibilants, as they're called by studio nerds like me). Moreover, the general preference for bright vocal sounds in modern productions only exacerbates the situation, because any static EQ will be a nonstarter. "If you use straight EQ'ing to take away your esses," says Manny Marroquin, "you take all the life and presence out of the vocal."[3]

Boosting the high frequencies of the main vocal compressor's side chain can help in this instance by encouraging the compressor to reduce the gain of high frequencies more stiffly. However, if sibilance is a big concern, then it's usually more effective to use a second compressor specially configured to combat it—the attack and release times that best suit the dynamic-range properties of your vocal performance may not be nimble enough to contain fast-moving high-frequency sibilants. Bob Clearmountain explains a particularly ingenious manual side-chain implementation that also finds favor with Jon Gass[4] and Manny Marroquin[5]: "I mult the vocal to the channel next to it, on which I'll roll out all the midrange and boost around 7kHz as far as it will go. I'll then send the output of the fader into its own buss and select a compressor on that channel to pick up on that buss, so basically it becomes a sort of loop in the channel. I put the compressor on fast attack and fast release, and link it to the original vocal channel, on which I turn on the dynamics. It works like a side-chain on the compressor and has the effect of a fast high-frequency limiter, its

threshold being controlled by the fader on the parallel channel. It's really effective, very fast, and it de-esses very smoothly and much better than any other de-esser I have ever used. Plus you can automate the threshold."[6]

Whichever setup you choose, it's important to work out which side-chain frequencies you need to boost. The easiest way to do this is first to use EQ boost in the main signal path with the aim of making the "s" sounds as overblown as possible. Once you know at which frequencies the sibilance resides, you can transfer that EQ setting to the compressor's side chain instead and boost until the sibilants come down to an appropriate level in the mix. However, there's a danger here: if you soften the "s" sounds too much, the vocalists will sound like they're lisping, so be careful not to go overboard with your de-essing, otherwise the singers might think you're taking the pith.

If you soften the "s" sounds too much, the vocalists will sound like they're lisping, so be careful not to go overboard with your de-essing, otherwise the singers might think you're taking the pith.

Pumping and Breathing

The flip side of compressors being generally less sensitive to high frequencies is that they're more sensitive to low frequencies, and that results in another frequently encountered compression side effect, which is usually most obvious on full-range, transient-rich signals such as drums and piano. What happens is that a low-frequency transient hits the compressor, triggering heavy gain reduction of the entire frequency range and dipping the sustain tails of any higher-frequency events that preceded it. Once the transient has passed, the gain reduction resets, bringing the higher-frequency sustains back into the mix. This lurching of the high-end balance in response to compression triggered from low-frequency events is commonly referred to as "pumping" or "breathing." If you want to hear a commercial example of this effect, check out Stardust's "Music Sounds Better with You," where the whole track is being pumped by the kick drum hits—a standard dance-music effect, although taken to a bit of an extreme in this case!

In some situations pumping can be desirable, because it lends sounds an artificial sense of aggression and power. This is partly because it increases the average level of the processed track, but also because it simulates the way the compression-like safety mechanisms within our own hearing system affect our perception when they're trying to protect the delicate apparatus of the inner ear from damagingly high sound levels. Drum sounds in rock music regularly use compression pumping for these reasons, and many dance tracks create an extra illusion of power for the main kick drum in this way, as in our Stardust example. Naturally, pumping won't be at all appropriate in other situations, and a dose of low-end shelving cut in your compressor's side chain can work wonders in reducing it. Indeed, some celebrated analog compressor designs,

FIGURE 13.3
DDMF's NYCompressor is one of only a few compressors that include variable onboard side-chain equalization, a feature that allows it to avoid unwanted pumping artifacts that other designs struggle with.

particularly those renowned for processing full mixes, incorporate low-end contouring of this kind in their designs.

13.2 MULTIBAND DYNAMICS

Side-chain equalization can help a good deal with refining the action of dynamics processors, but it can't overcome the inherent limitation of full-band processing—namely, that it always applies the same gain reduction across the entire frequency range. Consider a drum loop where the kick drum's low frequencies are far too prominent in the balance, yet when you try to compress them back into line you hear unacceptable pumping artifacts on the other instruments. Side-chain EQ may reduce the pumping, but it'll also bring back the balance problem, whereas normal EQ in the main signal path will thin out all the other sounds that happen between the kick drum hits. One way to square this circle is by splitting the frequency spectrum of the track into a number of different ranges (or "bands") so that you can give each range its own independent dynamics processing—a setup generally referred to as "multiband dynamics." In the drum-loop example, a multiband approach allows us to compress just the low end of the loop, reducing the overprominent kick frequencies without affecting the low end of intervening sounds such as toms and snares, or indeed introducing unnatural undulations in the balance of the hi-hats, cymbals, and percussion higher up the spectrum.

The logic of applying different dynamics treatments to different frequency bands is most obvious in the event that different instruments within a single

recording occupy contrasting frequency ranges. Slightly more challenging conceptually is the idea of applying multiband processing to a single instrument, so quite a good trick for making sense of this factor is to think of the different frequency ranges of the processed track as if they were actually separate mini-instruments in their own right, each responsible for a different aspect of the composite sound. So, for example, you might visualize three frequency ranges of an acoustic guitar track as if they were an ensemble of three separate mini-instruments: the sound hole boom, the sound of the strings, and the pick noise. All three of these mini-instruments need to have their own level in the mix, yet each may need very different dynamics processing to stabilize its own notional fader. So low-frequency limiting might help to make the resonant contributions of the sound hole "instrument" more consistent; low-ratio compression in the midrange might enhance the warmth and sustain of the strings; whereas threshold-independent transient reduction at the high end might help soften unforgivingly spiky picking. Building a mix balance is about gaining control over all the different sonic components of every track, and the power of multiband processing is that it increases the precision with which you can target each of those individual components.

DO-IT-YOURSELF MULTIBAND PROCESSING

There are lots of all-in-one multiband dynamics processors to choose from, and one seems to be bundled with every DAW these days. However, though there are a few extremely sophisticated specimens on the market (such as Flux Audio's brain-meltingly complex Alchemist plug-in), there always seems to be something that each specific model can't manage. Maybe you need more bands, or you want a type of dynamic processing that isn't provided, or you fancy experimenting with different crossover slopes. One way to get the facilities you feel you're missing is to build your own multiband setup from scratch, and although this might seem a bit hardcore, it's fairly straightforward to do by combining filters and full-band dynamics processors. The main hazard to avoid is that filters and dynamics plug-ins can both mess with the phase response of audio passing through them, so be on the lookout for any undesirable hollowing out of the sound from phase cancellation. "All the while I'm checking phase," says Tony Maserati, who regularly uses his own do-it-yourself band-splitting setups. "That's the most important thing whenever you're combining two of the same signal that's been EQ'd and compressed differently."[7]

Drum loops and acoustic guitars are good examples of sounds that frequently require the attention of multiband processing in the small studio, but there are lots of other occasions where it can also be warranted:

- It's common for bass guitar recordings made in small studios to be uneven at the low end, because of performance difficulties, a poor-quality instrument, or resonance problems in the cabinet or recording room. Upright bass can suffer too, although body resonances take the place of cabinet

resonances in this case. Either way, low-frequency compression can ride to the rescue.

- Kick drums that rumble on for too long at the low end can cause big problems in an up-tempo mix or in arrangements that demand good bass-line definition. Expansion of the drum's low-frequency range can tighten up the low end of your production considerably, improving clarity and punch.
- Synth lines with high-resonance filter sweeps can be a nightmare to balance if the filter's frequency peak is too prominent. You can't use normal compression, because the peak is there all the time, and EQ won't work either because the peak frequency doesn't stay still. If you compress the frequency range over several frequency bands, gain reduction will only occur in the band that the peak inhabits.
- Background synth textures (such as pads) and other atmospheric noises can encroach too much on the foreground of a mix if they include any transients. However, applying enough transient reduction, even with a decent threshold-independent processor, can involve unacceptably jerky gain-change artifacts in the otherwise fairly continuous backdrop. Splitting the transient processing across several bands gives you more transient reduction before side effects become distracting. I find this especially helpful when layering vinyl noise into a production, as many vinyl-noise samples include pops and crackles alongside the more steady-state hiss signal.
- Vocal sibilance problems can be targeted with multiband dynamics too, by compressing the high end. The trick is first to set a threshold that results in gain reduction only for the overloud sibilants and then to finesse the compression ratio so that you get the required gain reduction.
- All directional microphones inherently boost the recorded bass levels of sounds that are close to them (a phenomenon known as the "proximity effect"), so if you record any performer who moves backward and forward with relation to the mic, you'll get an uneven low end. Compressing the low frequencies can help salvage a more dependable level, but be prepared to spend time finding the best attack/release times and experimenting with the exact crossover frequency/slope to get the smoothest results.

The biggest single pitfall to avoid with multiband processing is the temptation to process across loads of different bands just because you can. In practice, the majority of mixdown scenarios that benefit from multiband treatment only require dynamics processing on one section of the frequency response. The previous chapters in this book should already have armed you with the understanding to avoid this trap, but it bears repeating that blanket multiband processing seldom makes any sense at all from the perspective of balancing a mix. In fact, one of the big reasons why the presets on all-in-one multiband processors are largely useless for mixdown purposes is that their approach is too indiscriminate—if you finally locate a preset that deals with your specific balance problem, you'll usually discover inappropriate settings wreaking havoc elsewhere in the frequency spectrum.

DYNAMIC NOISE REDUCTION

Unwanted background noise has vexed small-studio users since time immemorial, but digital technology now gives us many more options for fixing it in the mix. The simplest remedy is gating, but the disadvantage of this is that it only reduces noise levels in the gaps between instrument/vocal sounds. A slightly more sophisticated approach involves a kind of dynamic low-pass filter, the cut-off frequency of which tracks the signal level. Although this isn't that much better than a straight gate, it does work better at reducing noise during note sustains, which are typically less bright and therefore mask hiss less. Multiband expansion can give you another slight increase in the efficiency of hiss reduction, but it's not going to help a great deal with separating desirable high frequencies from the noise floor.

Unless, that is, you can increase the number of expansion bands well beyond the four or five offered by most all-in-one multiband processing plug-ins, in which case you can actually gain a surprising degree of noise reduction with little apparent loss of desired signal. However, controlling hundreds of bands of expansion would be like juggling a family of slimy amphibians (the Osbournes, say), so dedicated noise-reduction processors use a different control method, which semi-automatically configures the controls in response to an isolated sample of the unwanted noise. As such, there's not a tremendous amount to say about using dedicated noise reduction, beyond reiterating that it's usually best tackled from the first slot in any track's processing chain. What can be informative, though, when you first use one of these beasts is to turn the noise reduction up way too far so that you can hear what the processing artifacts sound like. Then when you're trying to gauge the depth of the proper processing, you'll be more keenly aware of what to listen out for. If you don't have multiband noise reduction on your DAW system already, check out Voxengo's affordable cross-platform Redunoise plug-in.

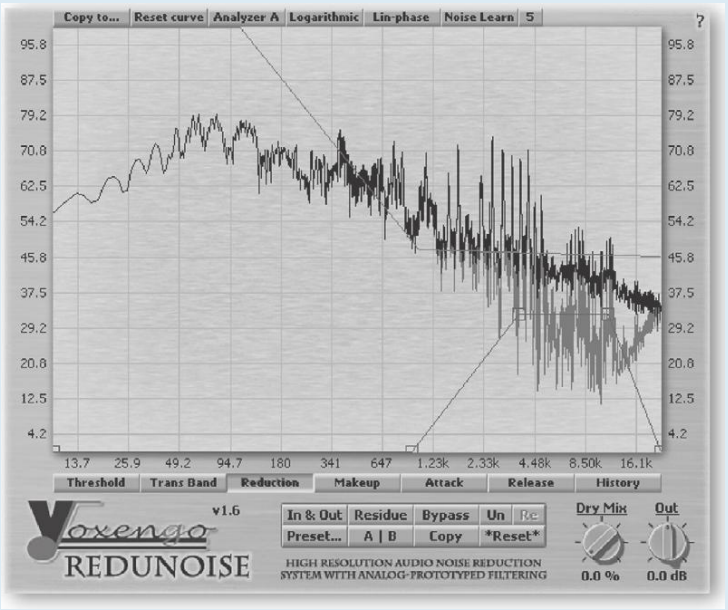

FIGURE 13.4
An affordable option for dynamic multiband noise reduction is Voxengo's cross-platform Redunoise plug-in.

I also want to stress that the process of setting up and evaluating the action of each dynamic processor within a multiband plug-in should involve exactly the same techniques and thought processes that pertain when using a single full-band processor. Again, this is where it helps to think of each frequency band as if it were a separate track, because that way you're more likely to give it the amount of attention it deserves. If you expect to get results equally quickly from full-band and multiband dynamics plug-ins, then you'll cut too many corners. Attack and release times in particular need to be reassessed for every frequency band, because lots of multiband devices seem to give you identical attack/release values for all bands by default, regardless of the fact that signals in different frequency ranges move at very different speeds. With multiband compressors, you should also be wary of automatic gain compensation, because some models don't allow you to switch it off. If you can counteract any automatic gain boosts with manual makeup gain cuts, you'll be at less risk of playing psychological tricks with yourself due to loudness changes.

Multiband dynamics can be used for parallel processing, just as full-band models can, and this opens the doors to a wealth of natural-sounding frequency enhancement effects, especially when using compression.

When it comes to deciding on the crossover frequencies between processing bands, my first recommendation is to use only as many bands as you really need, because this keeps any side effects from the crossover filters to a minimum. If you only need to process the high end of a track, then two bands will be fine, regardless of how many your specific plug-in offers by default or can provide in theory. The exact crossover frequencies will, of course, be critical to gaining the best control over the spectral regions you're concerned with, and it can help here to solo each band in turn to refine these settings, although a decent spectrum analyzer may help too. Some plug-ins offer a choice of slopes for their crossover filters, which is actually a useful option, because sharper slopes tend to work better for narrow-band troubleshooting (such as controlling wayward pitched resonances), whereas gentler slopes normally sound more natural for broadband enhancements (such as increasing the high-frequency sustain of a piano).

One final tip: multiband dynamics units can be used for parallel processing, just as full-band models can, and this opens the doors to a wealth of natural-sounding frequency enhancement effects, especially when using compression. Just try to choose a multiband plug-in that doesn't spoil the party by messing with the processed signal's phase response and introducing comb-filtering nasties.

13.3 DYNAMIC EQUALIZERS

An alternative way to achieve frequency-selective dynamics is to swap the static gain controls of a regular equalizer for gain-control elements that can respond

to the input signal in real time, thus creating what's known as a dynamic equalizer. Dynamic EQ is functionally very similar to multiband dynamics, inasmuch as both adjust the time-domain characteristics of localized areas of the frequency spectrum, so you'll find that many jobs can be accomplished just as readily with either type of treatment. However, dynamic EQ has a few tricks up its sleeve that go beyond the scope of most multiband dynamics processors, so it warrants further investigation.

A dynamic equalizer works by making each filter's gain control act in the same way as the gain-change element in a dynamics processor. Depending on the model, there may be the option for it to mimic a compressor, an expander, or even a threshold-independent transient enhancer. For each band you therefore need to decide on the filter type, frequency, and bandwidth and then set up the band's dynamics processing type and controls (Threshold, Makeup Gain, Ratio, Attack, Release, and whatever else) to govern how gain is applied. If that sounds like a lot of work, you're not wrong, but the upside is that once you've got the hang of dynamic EQ, it can bail you out of situations that practically nothing else can, by virtue of the fact that it can operate on precise frequency regions—or even single frequencies in some cases.

FIGURE 13.5
Dynamic equalizers, such as 112dB's Redline Equalizer (*top*) or Platinumears IQ4 (*bottom*), can defeat some of the most dastardly mix gremlins, but they do take some work to set up properly.

Lead vocal parts are where dynamic EQ probably sees most use. Sibilance reduction is one of the options on the menu, but although a dynamic EQ can be exceptionally good at that function, the real star turn of this processor is its ability to contain sporadic narrow-bandwidth frequency peaks. The reason this is so useful with vocals is that a singer creates different timbres and vowel sounds by physically altering the resonant qualities of the throat, nose, and mouth. With some singers, you'll find that certain combinations of note pitch, delivery force, and vowel sound can briefly create an enormous resonant peak somewhere in the frequency response (normally somewhere above 1kHz), and this makes it almost impossible to balance that vocal section satisfactorily. When the rest of the vocal sound feels in balance, the resonance pokes you in the ear with a pointy stick whenever it surfaces; if the fader is set to balance the resonance, then the rest of the vocal disappears into the mix. Normal dynamics don't help, because the resonance doesn't add substantially to the vocal's overall level, and normal EQ doesn't help either, because any equalization savage enough to zap the resonance will also administer the last rites to the overall vocal timbre. Mick Glossop is one of several producer/engineers who have touched on this problem in interview: "It's very common for a singer's voice to develop a hard edge when he's singing loudly, which doesn't happen when he's singing quietly…. The problem with using static EQ to get rid of the hardness is that it might solve the problem of the vocal in the chorus, but when the quiet verse is being performed, it will lack presence, because you've taken out some 4 to 5kHz."[8]

DEDICATED DE-ESSERS

A de-esser is a specialized processor designed to reduce sibilance levels in vocal parts. I've already mentioned how you can implement de-essing from first principles using a compressor (in conjunction with side-chain EQ), a multiband compressor, or a dynamic equalizer, and all of these methods can give good results. Almost all dedicated de-essers merely provide a more streamlined version of one of these methods, with an interface that's quicker and more convenient to use, but there is also another much less common type of de-esser that operates in a different way, chopping out sibilant time regions from the audio and allowing you to process them separately. This last design can give you even greater control over the processing of sibilant regions, and although in my opinion it verges on overkill for everyday tasks, you may wish to investigate it if sibilance is proving truly overwhelming in your mix. The main point is that different de-essers respond differently, so if you've got sibilance worries that cause one algorithm to break out in a sweat, then do a shootout of some other brands.

I usually put any de-essing right at the end of the processing chain on vocal tracks, because otherwise it doesn't counteract any sibilance increase from down-the-line compression or EQ plug-ins. However, there's one important exception to this general rule: I'll always put it before any type of high-frequency distortion enhancement, such as the Aphex Aural Exciter. This is because if the enhancer comes first, then the high-level sibilance will create lots of distortion harmonics, adding enough harshness to grill your

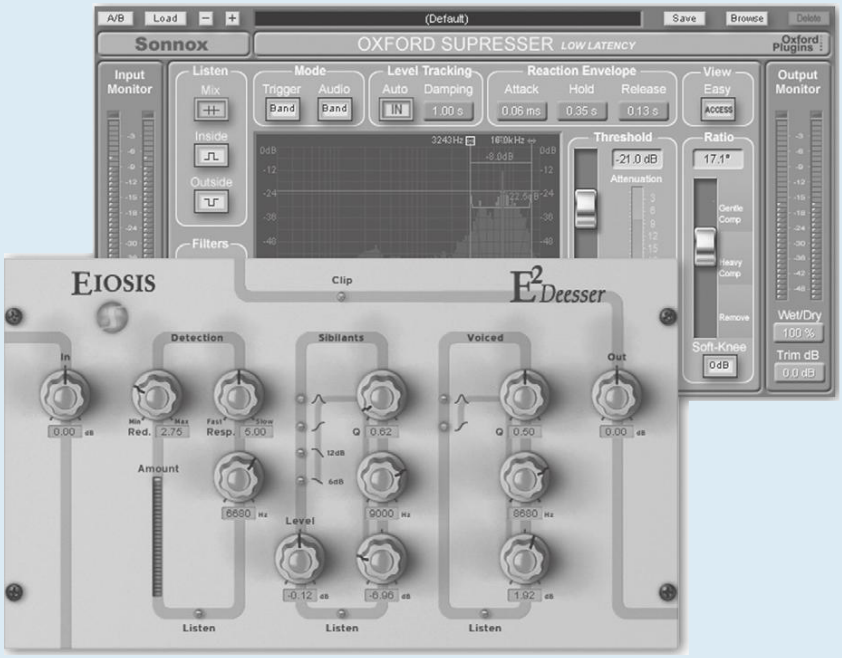

FIGURE 13.6
Although there are some powerful and elegant state-of-the-art de-essing plug-ins available, including Sonnox Oxford Supresser (*top*) and Eiosis E2 (*bottom*), you can also achieve excellent de-essing using less esoteric processing.

eardrums to medium rare! Then even if the de-esser reduces the level of the sibilant regions, it won't soften their tone, so you can get into a situation where the de-essing starts making the vocal lisp well before the sibilants are actually sounding polite enough. De-essing before the enhancer means that the sibilants drive the distortion less hard, so they maintain a smoother sound.

Although de-essers are clearly designed to do what they say on the box, you can often find other uses for them if you understand what they're doing behind the scenes. Most available de-essers work by compressing the high frequencies, so you can often use them to reduce other high-frequency undesirables too, such as an acoustic guitar's fret noise.

With dynamic EQ, however, it's fairly straightforward to use a narrow peaking filter in compressor mode to control resonant peaks like these; the tricky bit is diagnosing the malaise in the first place and knowing exactly where the resonance lives so you can target it accurately. In fact, this particular balance problem is one of the slipperiest customers you're likely to encounter when mixing, because it takes a while to learn to recognize the specific piercing harshness

(often on notes where the singer is stretching beyond her comfort zone in terms of pitch or volume) that is normally the first subjective clue to its existence. The second tip-off I typically get is that ordinary EQ draws a blank in terms of tackling the harshness effectively, and at that point I usually zoom in, loop one of the harsh-sounding vocal notes, pull out my highest-resolution spectrum analyzer, and scour the display for any spectral peaks that are shooting into orbit. Sometimes there's only one to deal with, but I've had situations where there are four or five hopping up out of the mix at different points in the track, and each had to be carefully dealt with before the vocal would sit properly in the mix.

Few other instruments present this particular problem because they don't change the way they resonate, but there are nonetheless still a few other instances where dynamic EQ can save the day. The resonant sweep of a wah-wah guitar part might periodically coincide with a strong cabinet resonance in your guitar combo, for instance; the fundamental frequency of a certain upright bass note might vary considerably in level depending on whether that note has been stopped or played as an open string; or your pianist might be too heavy with his right thumb on one note of an important riff. Dynamic EQ may be fiddly, but it can also be your last resort when the going gets tough. If you don't have access to such processing on your own DAW system, then there's a good freeware VST-format plug-in from Platinumears called IQ4, which is well worth an audition.

CUT TO THE CHASE

- Frequency-selective dynamics processing allows you to deal with balance problems that occur in both the time domain and the frequency domain at once and can therefore overcome some of the trickiest mixing obstacles. Despite the apparent complexity of such processing, using it is mostly just an extension of what we've already covered in previous chapters.
- Any normal dynamics processor can be made frequency-selective either by using it in a parallel setup with an equalized return channel or by equalizing the input to the processor's level-detection side chain. Side-chain equalization is particularly useful for refining a gate's triggering, as well as for reducing the sibilance emphasis and pumping/breathing side effects of heavy compression.
- If you treat separate frequency regions within a given track as if they were separate tracks in their own right, the logical application of multiband dynamics processors becomes easier to grasp. Some tips: steer clear of presets; avoid automatic gain-compensation routines; spend enough time refining crossover settings; and don't feel you have to use every available band, or indeed all the processing available in each band.
- Dynamic equalizers can both mimic many of the effects of multiband dynamics and go further by dealing with dynamics problems that have very narrow bandwidth.

- De-essers can operate in many ways, so if the first plug-in you try doesn't deliver the goods, make a point of auditioning some alternatives. De-essing normally works best at the end of a track's processing chain, although it's best to apply it before any distortion-based high-frequency enhancements.
- If you have to deal with a lot of noisy tracks, then you might need to call on dedicated multiband noise-reduction processing.

Assignment

- Check what de-essing and multiband dynamics plug-ins are available on your DAW system, and if the choice appears limited, then consider investing in third-party additions.
- Mute all the tracks in your mix and rebuild the balance, experimenting with frequency-selective dynamics processing wherever you feel it may help solidify the balance.
- Make a note of any faders that still feel unstable.

www.cambridge-mt.com/ms-ch13.htm

This page intentionally left blank

CHAPTER 14

The Power of Side Chains

I introduced the idea of dynamics side chains in Chapter 13 but have so far assumed that the level-detection signal path is fed from the same source as the main processing path—albeit perhaps with different equalization. As it turns out, there are sometimes very good reasons for feeding a dynamics processor's side chain from a different track entirely, and many dynamics plug-ins (most commonly full-band compressors, expanders, and gates) are now capable of accepting an external side-chain input.

The most basic reason for using an external side-chain signal is if you can't get the gain reduction to respond correctly based on the signal levels of the track you're processing. So, for example, imagine you have two mics on your rack tom: one on top and the other poked inside the drum from underneath. The top mic may well give the better sound, but it may also prove impossible to gate reliably because of snare spill. By using the inside mic's signal to "key" (via the side chain) the gate on the outside mic, you can achieve much better triggering, without any need to use the dodgy sound of the inside mic in the balance at all.

Side chains are about much more than this kind of troubleshooting, though, because they enable tracks to interact with each other to improve the overall mix balance. Probably the most commonplace application of this capability is in rock music where overdriven electric guitars are masking the lead vocal's high frequencies. Although you may be able to carve out space for the vocals with EQ in this scenario, the danger is that you'll simultaneously make the rest of the arrangement sound like *The Care Bears Movie*. So to take some of the pressure off the EQ processing, you can set up compression on the electric guitar parts that responds to a side-chain input from the vocal track. Whenever the vocalist sings, that triggers gain reduction on the guitars and therefore reduces the unwanted masking; but when the vocal is silent, the guitars return to their original level in the balance. Just a decibel or two of gain reduction can work wonders, although you do have to be careful that you don't go so far that the gain changes become distracting.

Mixing Secrets for the Small Studio.

There are plenty of other situations where similar dodges can be of assistance:

- To get rid of a nasty snare-spill component on your hi-hat mic, use the snare close-mic to key a hi-hat compressor.
- To achieve the tonal benefits of strong drum ambience for individual instruments without making a mess of the mix, try keyed gating on your room mics. The most famous example of this, of course, was Hugh Padgham's trend-setting drum sound for Phil Collins's "In the Air Tonight," but before you write it off as just some cheesy 1980s gimmick, bear in mind that Chris Thomas and Bill Price also used the idea for the Sex Pistols' debut album *Never Mind the Bollocks*: "[It] involved me keying different ambience mics off the drums as they were being hit," recalls Price. "Chris Thomas's suggestion that we could shorten the ambience with gates, providing more without it sounding too distant, all made total sense to me."[1]
- To boost the snare sound in your overhead mics, use the snare close-mic to key a limited-range gate on the overhead channels and then fade them up to rebalance the cymbal levels. (Another time-honored trick, this one.)
- To avoid excessive low-end buildup, use the kick drum part to key a compressor on the bass. (Some people prefer to route the kick and bass through a single compressor to achieve a broadly similar result, but I don't find that method quite as controllable myself.) "I often do this when the kick is being stepped on by the bass," says Jason Goldstein. "So every time the kick hits, the bass ducks 2dB or so just for a moment. When you have a [prominent bass], you can't do too much."[2]
- To give a rhythm part or pad sound more rhythmic impetus, insert a limited-range gate on that channel and key it from an in-tempo metronome click or a simple repetitive audio track (perhaps created specially for the purpose). Work with the gate's attack and release times to get the most musical pulse.
- To add low-end welly to a kick drum, route a constant low-frequency sine-wave tone through a gate, and then trigger it from the drum sound. This was a much-loved trick in the disco era and can still pay dividends nowadays in dance and urban styles. It doesn't really matter whether you create the tone using a simple test oscillator or a more fully featured synthesizer, but I'd suggest starting with a frequency of around 50Hz and then tweaking that by ear to find the best fit for your particular mix. If it's out of tune with your bass line, for instance, it can really destabilize the musical harmony.
- To prevent long echo and reverb effects from swamping a lead vocal, key a compressor on each effect return from the lead vocal signal. You'll be surprised how much gain reduction you can get away with before anyone starts feeling seasick. "Dynamic delays… will duck out of the way, coming up only at the tail end of words and phrases," says Goldstein. "I almost never use a straight delay."[3]

DUCKING

Although a lot of engineers use keyed compressors to dip the levels of backing instruments in response to lead vocal levels (as discussed in the main text), I don't personally like this approach too much. In the guitars-versus-vocals case, for example, the quietest vocal notes will cause the least reduction in guitar levels, even though the quietest notes are the ones that need the most help. My preference, therefore, is to use a dedicated ducker instead. A ducker operates exactly like a gate, except that it opens when a gate would close, and vice versa—if you can already set up a gate, then operating a ducker should be a no-brainer. What's different about using a ducker for this kind of keyed mix-balancing task is that it will always introduce the same gain reduction for all the vocal notes, and though that's still not perfect, to my ears it's nevertheless a significant improvement over keyed compression.

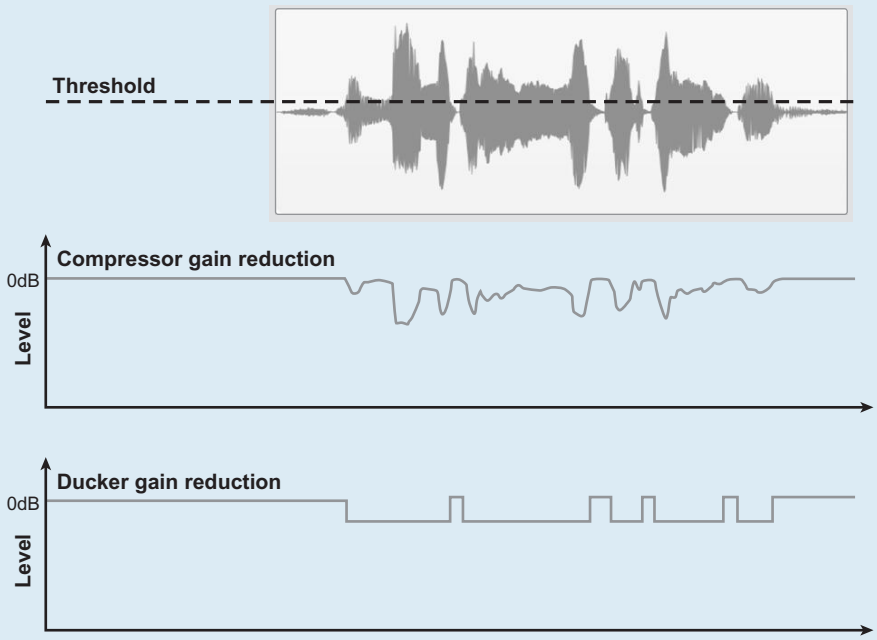

FIGURE 14.1
Comparing the gain-reduction actions of a compressor and a ducker. Notice that the compressor reduces the gain most when the vocal is highest in level, which isn't ideal for the purposes of side chain-triggered rebalancing of other competing tracks in the mix.

A stumbling block for DAW users, though, is that there are few proper side chain–enabled duckers available. Fortunately, however, there's a little workaround you can do to arrive at an identical effect in pretty much any DAW. Here's how it works within the context of the vocal/guitars scenario. Set up a gate as a send effect, send to the gate

(*Continued*)

from the guitar channels, feed the gate's side chain from the vocal channel, and then invert the polarity of the gated return. What should happen with this setup is that every time the vocal is present, the gate will open to let through a burst of inverted-polarity guitar signal, and this will silence the main guitar channels in the mix by virtue of phase cancellation (assuming that the gate introduces no latency delay). All you now need to do is decide what level to set the fader of the return channel: the lower you have it, the less severe the phase cancellation, and the finer the ducking—as the Reverend Spooner might have put it.

This setup has a nifty application in dance music, too: it's a particularly elegant way of creating kick-triggered gain pumping across multiple tracks, because you can determine the amount of ducking for each instrument by adjusting its send level to the parallel gate. So if you wanted severe ducking of your synth pads and echo/reverb effects, you'd turn up the gate sends on those channels, while lead vocals or synth leads might have lower send levels to avoid the pumping making a nonsense of the main hooks. However, my favorite ninja ducking technique of the lot is when you add a linear-phase high-pass filter into the gate's return channel and trigger the gate from your lead vocal. Now any track that sends to the gate will be ducked whenever the singer is present in the mix, but only at high frequencies, and only to a depth determined by that track's send-level control. Now *that's* mixing power!

14.1 KEYED MULTIBAND DYNAMICS

Once you get the idea of side-chain triggering for full-band dynamics, you don't have to extrapolate too much further to conceive the idea of keyed multiband dynamics, although you're unlikely to come across a balance problem that demands this kind of firepower more than once in a blue moon. The most common situation would probably be the vocals-versus-guitars situation I've already mentioned, because the multiband strategy would enable you to apply gain reduction to the guitar parts primarily in those spectral regions that most heavily mask the vocal sound. Similarly, in the kick-plus-bass example presented earlier, you might retain a better subjective balance for the upper frequencies of the bass if the kick-triggered compression were restricted to low frequencies.

A practical hitch, though, is that only very few real-world multiband dynamics processors offer external side-chain access, so if you want to explore such avenues, you may have to create your own do-it-yourself multiband setup by combining filters and keyed full-band processors. In some cases an alternative might be to EQ the return of a keyed parallel dynamics setup—perhaps isolating the midrange frequencies of the overheads in the snare-plus-overheads example described earlier, so that the gating boosts the snare more than the cymbals in the overall balance.

CUT TO THE CHASE

- Some mix imbalances can only be resolved if dynamics processing on one channel is controlled by a signal from another channel. Full-band dynamics will be most straightforward to use in this way, because few all-in-one multiband processors offer external access to their level-detection side chains, but there is little to stop you from creating keyed frequency-selective dynamics using equalized parallel processors or do-it-yourself multiband setups should the need arise.
- Duckers are often better suited to keyed mix-balancing tasks than compressors. However, if you can't get hold of one for your system, a keyed parallel gate channel can be polarity inverted to arrive at an identical ducking effect—a setup that has additional power-user applications.

Assignment

- Investigate the side-chain routing facilities within your own DAW system, and make a note of which plug-ins will accept an external side-chain feed.
- Mute all the tracks in your mix and rebuild the balance again, experimenting with keyed dynamics processing to see if it can stabilize the remaining faders.

www.cambridge-mt.com/ms-ch14.htm

This page intentionally left blank

CHAPTER 15

Toward Fluent Balancing

You should now have enough tools and techniques at your disposal to create a good static mix balance. If you've worked through the end-of-chapter Assignment suggestions, you should also have gained some practical insight into which processes suit different balance problems within your musical style of choice. In other words, you'll have started to interpret the subtle whisperings of fader instability, which are trying to tell you what the music needs. As I openly admitted back at the start of Chapter 8, the step-by-step procedure implied by the Assignments has no foundation in real-world mixing practice, but I make no apologies for that, because it makes the vastness of the subject easier to grasp initially. However, now that my overview of balancing procedures has run its course, it's high time we removed the didactic scaffolding and examined how all these methods fit into the more fluid workflow of the professionals.

The key to fluent balancing is to let the faders set the agenda, so rather than rebuilding the balance multiple times and adding different processes with each rebuild, you only actually need to build the balance once. As you introduce each new track into the mix, you address its panning, filtering, and phase/polarity, and then try to set a balance. If you can't find a stable fader setting, then you keep experimenting with different processing until you can. Whether it takes a single subtle equalization tweak or unspeakable abuse at the hands of a whole gang of bloodthirsty plug-ins, you keep going until you get a solid balance without unacceptable processing artifacts. The goal is to get every aspect of the track to balance: every pitch and every noise; every transient and every sustain; every moment in time and every region of the frequency spectrum. When you've succeeded (or else exhausted every processing option you have available), you move on to the next track and start again.

> When push comes to shove, the needs of the balance should always be your first priority—if your mix doesn't have balance, no one's going to give a monkey's whether any individual instrument sounds stunning.

Of course, any given track may also feel subjectively unappealing, even at an appropriate balance, in which case you may also need to experiment with a

certain amount of "suck it and see" creative tonal shaping. However, this kind of mix treatment is applied in a very different manner, so it has to be evaluated on its own terms ("Do I like the sound better now?") and shouldn't be confused with balance processing ("Can I hear everything properly?"). Moreover, when push comes to shove, the needs of the balance should always be your first priority—if your mix doesn't have balance, no one's going to give a monkey's whether any individual instrument sounds stunning.

Although it's vital to have a well-developed plan of attack and to try your best to deal with each track's balance before adding in the next, it's also important to recognize that real-world mixing can never be a truly linear process, so no one's going to shoot you for tweaking the processing of your first track while balancing your twenty-seventh. On the contrary, it's often only when you're desperately trying to shoehorn your final tracks into the few remaining mix crannies that you get an accurate picture of what aspects of the more important tracks are truly indispensable. So if you become aware of a previously unnoticed fader instability on any earlier track, then go back and deal with it before introducing any further instruments. (And if you can manage to make this look completely erratic to the uninformed observer, you might also develop an aura of professional mystique into the bargain!)

SMALL IS BEAUTIFUL

There is a common misconception among small-studio mix engineers that it takes big, bold processing moves to bring about large sonic changes at mixdown, which is why so many amateur mixes suffer from overprocessing. However, if my experience with dozens of "Mix Rescue" remixes has taught me anything, it's that the most dramatic mix transformations are usually brought about primarily through lots of small mix tweaks. It doesn't matter if the benefit of any individual plug-in setting seems almost insignificant on its own, as long as its benefit outweighs any unwanted side effects, because even miniscule improvements add up surprisingly swiftly. So think like a sculptor: you won't get too far by clouting your piece of stone a couple of times with a sledgehammer; you need to chip away at it bit by bit to get anything worthwhile.

Here's another point to be absolutely clear about: although fader instabilities can give you a hundred clues, all processing is one long experiment. "Real music is about experimentation," to quote DJ Premier.[1] You simply cannot know that a given process will work before you try it, because every mix is different. Only your ears can decide if any tone/balance improvement you've dialed in is worth the price you've paid for it in terms of processing side effects. So it's normal to make duff calls when it comes to choosing the right processor, and it's all right to strip away anything that isn't working and start afresh—even if that means zeroing absolutely everything and beginning the balance again from scratch. "I will often restart mixes three or four times," reveals

Fabian Marasciullo. "Put everything back to zero and try again, reblend and EQ everything, and put the vocals back in."[2] Every time you flunk out, you'll have eliminated a setting that doesn't work, and that takes you a positive step closer to finding a setting that does. "I think a lot of the thing with music is pilot hours," muses Youth. "To go through all your mistakes and know not to go up those cul-de-sacs and fall into those elephant traps."[3]

And finally, low budget is no excuse for bad balance. If you can't get a balance with the bundled plug-ins in any mainstream DAW, then it's not the fault of the gear. A big budget will only make mixing quicker and more pleasurable if you can already get a great mix with no-frills gear. But don't take my word for it—ask Frank Filipetti, who won his 1998 Best Engineering Grammy for *Hourglass*, an album recorded in James Taylor's house using a first-generation Yamaha 02R and three Tascam DA88s. "In the end, your ears, your mind, your musical abilities are what it's all about. Put a George Massenburg, a Hugh Padgham, a Kevin Killen together with any kind of gear, and you'll get a great-sounding record.... If you have a sound in your head and you know what it is, you can get it. You may have to work harder on lower-quality gear, but you can get it."[4] Tony Visconti agrees: "I've heard people make very bad records on expensive gear. The gear does not dictate the quality. It's how you use it."[5]

You don't need every esoteric specialist processor to get the job done either, because there are so many workarounds when it comes to mix processing. If you don't have a de-esser, for example, then try a compressor with EQ in the side chain, or high-frequency compression from a multiband device or do-it-yourself equivalent, or an 8 kHz peaking cut from a dynamic EQ in compressor mode, or a polarity-inverted parallel gate with its side-chain filters set to isolate the sibilant frequencies, or a static EQ under the control of your DAW's mixer automation system. Or just mult the sibilant audio sections to another track and set their fader level separately! The gear is not the issue; it's you, your ears, and your monitoring that make a balance work.

CUT TO THE CHASE

- Fluent balancing technique is about reading the instabilities of your faders as they arise and drawing on the relevant tools in your processing arsenal as required. There is room within this process for subjectively enhancing sounds, as long as this doesn't undermine the success of the overall balance.
- Don't expect the mixing process to work in an entirely linear way, because it's normal to reassess early processing decisions later in the mix, once mix real estate begins to become more scarce.
- All processing is experimentation. If any processing experiment doesn't work, then don't think twice about ditching it and starting over—you'll still have learned something useful about the mix problem you're grappling with and will be better equipped to solve it on the next attempt.
- Low-budget processing is no excuse for delivering a bad balance.

Assignment

- Save your mix as it currently stands, and then open a copy of it, resetting all the channel settings and stripping out all the plug-ins so that you're left with a completely blank slate. Now redo the whole balance in one go, introducing the tracks as before, reading each fader in turn, and applying whatever processing is required to nail down a good static balance.
- Once you've completed and saved your new balance, compare it with your first attempt. As likely as not, the second version will have been quicker and easier to do, because you didn't go down as many blind alleys, and it'll probably also sound better into the bargain.

www.cambridge-mt.com/ms-ch15.htm

PART 4
Sweetening to Taste

Once you've achieved a good balance, you're effectively on the home straight. Even if you do nothing else to the mix, your listener will at least hear everything on the multitrack clearly and at an appropriate level, and that should guarantee that the quality of the music comes across. Nonetheless, there are lots of bonus points to be gained in most productions by adding further "sweetening" effects to an already creditable balance, the aim being to present the recorded material in a more flattering light. As Lee DeCarlo explains, "Effects are makeup. It's cosmetic surgery. I can take a very great song, by a very great band, and mix it with no effects on it at all and it'll sound good, but I can take the same song and mix it with effects and it'll sound fantastic! That's what effects are for."[1] It's these additions that I want to focus on in the remaining chapters, as well as discussing the final stages a mix goes through before being signed off as finished product.

Once you've achieved a good balance, you're effectively on the home straight. Even if you do nothing else, your listener will at least hear everything clearly and at an appropriate level, and that should guarantee that the quality of the music comes across.

If you're wondering why we're only dealing with sweetening effects now, then think of this as an extension of the idea of building the balance in order of importance. Such effects are, pretty much by definition, less important than the tracks that they're designed to enhance, so it stands to reason that you should add

them only once your raw tracks are balanced. That way you know how much space there is in the mix to accommodate the sweeteners, and are less likely to obscure important musical details by over-sugaring. Russ Elevado says, "I usually will not turn on any reverbs until halfway through the mix.... I try and achieve a big sound without resorting to reverb. So I just keep chipping away until the track is pumping. Then once I'm happy, I start to think about what I might want some room or ambience on."[2]

CHAPTER 16

Mixing with Reverb

The most widely used sweetening effect in record production has got to be artificial reverberation (or reverb for short), an effect that can generate complex patterns of echoes in response to an input signal. Artificial reverberators were originally designed to simulate the sonic reflections you get from boundaries in a real acoustic space, allowing you to add realism to unnatural-sounding close-miked recordings. However, reverb has a much broader brief in record production these days, and artificial reverberators are now used as much for creative purposes. Indeed, many classic reverberation devices don't produce anything like the sound of a real room, but that hasn't diminished their enduring popularity!

16.1 FIVE ENHANCEMENTS AT ONCE

Part of the reason why reverb is incredibly useful at mixdown is that it can enhance several aspects of your sonics simultaneously. As I see it, it has the power to enhance the following significant elements:

- *Blend.* Reverb can increase the extent to which any individual track blends with the rest of the production, making disconnected instruments sound as if they belong together and giving the mix as a whole a more cohesive sound. An instrument that isn't at all blended sounds upfront and close to the listener, whereas an instrument that blends well is sucked more into the background, away from the listener. Therefore, the idea of blend is also closely related to the idea of front-back "depth" in a mix: less blend brings a track toward you, whereas more blend pushes it away from you.
- *Size.* Artificial reverb can increase the apparent dimensions of your mix's acoustic environment, making it sound as if your tracks were recorded in a larger (and maybe better-sounding) room than they actually were—which is one way to make low-budget projects appear more expensively produced. In addition, if any given instrument excites the simulated reverberation particularly audibly, then it creates the illusion that this instrument is large and powerful in itself, even if it's low in level and well-blended with the overall balance. So, in a nutshell, reverb can increase the size both of the whole mix and of individual instruments.

Mixing Secrets for the Small Studio.

- *Tone.* The echoes that make up a reverberation effect have the capacity to phase-cancel with the dry track once they are added into the mix. The resulting comb filtering will alter the subjective tone of the instrument. Irregularities in tonal quality of each of the echoes will also contribute to this effect.
- *Sustain.* Because echoes are essentially delayed versions of the effect's input, any reverberation effectively increases the sustain of a dry sound it's added to. However, reverbs never sustain equally across all frequencies, so the tonal characteristics of the added sustain are a key parameter to be controlled.
- *Spread.* The simulated echoes in most artificial reverberators are distributed across much of the stereo image. This spreads information more evenly across the picture and may also increase the apparent stereo width of individual processed tracks, or indeed the whole mix.

Although reverb's multifaceted nature is good news in principle, the big catch is that a lot of reverbs tend to apply all these enhancements at once, whether or not you want them to! Furthermore, within a specific reverb effect, the amount and nature of each individual enhancement can be fiendishly difficult to adjust independently of the others. The result is that less experienced engineers usually find that they can't get enough of one enhancement without overdoing another: they can't get enough blend without muddying the mix tone, perhaps, or they can't enrich an instrument's sustain enough without apparently transporting it to the Taj Mahal!

Whether you fall foul of these mix traps has little to do with how much you know about the physics of natural reverb or how well you understand all the inscrutable-looking widgets on a typical reverb plug-in. The main secret to

FIGURE 16.1
Almost every reverb plug-in has a control somewhere that sets how much signal with effects and how much signal without effects it sends to its outputs. These screens show a selection of different configurations.

getting decent results swiftly is learning how to create reverb effects that provide only a subset of the possible enhancements. Once you have a set of more specialized effects at your fingertips, it's more straightforward to apply them in combination to achieve precisely the reverb sweetening your mix demands.

16.2 ESSENTIAL REVERB CONTROLS AND SETUP

To create these kinds of reverb effects, you need to understand a handful of essential truths about reverb plug-ins. The first thing to realize is that reverb is almost always best applied via a send-return effect configuration, so that a single effect can be accessed from every channel of your mix. For that to work properly, you need to check two factors:

- That your reverb plug-in is only outputting processed effect ("wet" signal) and not any unprocessed sound ("dry" or "direct" signal), otherwise sending to the effect from a channel will also alter that track's level in your finely poised mix balance. Some plug-ins may have a single Wet/Dry or Direct/Effect Mix control, whereas others might have independent level controls. Whatever the interface, though, make sure to turn the Dry/Direct signal all the way down or switch it off completely.
- That the individual sends that feed the reverb are taken from the channel signal path post-fader. This means that the balance of wet versus dry sound will remain constant if you adjust the send channel's main fader, and you won't have any ghostly reverb shadows remaining in the mix if you fade out the channel completely at any point.

FIGURE 16.2
The most important controls on a reverb are those that set how fast it decays, but different reverb plug-ins have different names for these, including Size, Damping, Decay, Time, and Length.

Because reverb is such an important studio effect, most DAWs tend to provide more than one algorithm, and numerous third-party plug-ins are available too, including lots of freeware and shareware. If you don't already have a preference for one or the other, then don't worry too much about which you choose to start with, because a lot of reverbs these days are pretty usable and it's often tricky to know how well a given reverb suits a particular task until you hear it in action. One little tip, though: good-sounding reverbs normally require a fair bit of CPU power, so be wary of any that are slimline in this regard, no matter how flashy their graphics look.

As far as all those reverb controls are concerned, the only ones that you absolutely have to know about are those that adjust how fast the echoes die away. On simple plug-ins there might be a single control for this, usually called Length, Decay Time, Reverb Time, or RT60. Where there isn't anything like that, then sometimes this parameter is adjusted by manhandling the decay "tail" on a graphical representation of the reverb envelope. Alternatively, you may find that Room Size and Damping are supplied instead, and although these fulfill a broadly similar function, they are also more likely to change other aspects of the reverb character

REVERB DESIGNS: A BRIEF FIELD GUIDE

You don't need to know the full history of how reverb effects were developed to use them effectively at mixdown. However, a little background knowledge about different reverb designs can nonetheless speed up choosing an appropriate plug-in or preset for each particular task, so here's a quick survey of the main options:

- *Chambers.* The earliest type of added reverb, created by sending signals to speakers in an unused room and then capturing the reflected sound with mics.
- *Plates and springs.* Two early types of electromechanical reverb, which were widely used during the 1960s and 1970s. The audio signal is used to set off vibrations in a bit of metal, and then the reflections of these vibrations are captured using pickups. Neither design is good at simulating realistic spaces, but both are nonetheless still highly regarded in the industry today, being well-suited to tonal and sustain enhancements.
- *Digital algorithmic processors.* A type of reverb that rose to prominence during the 1980s and had a huge impact on the sounds of that era. Because the reverb reflections are created using mathematical models, the user has a lot of control over them. Although some algorithmic processors are not very realistic (especially early models), a number of these are nonetheless prized for their unusual spatial and tonal qualities.
- *Digital convolution processors.* This recent technological development allows you to record the exact pattern of reflections from a real acoustic space as an "impulse response" file, which can then be used to recreate that reverb with unparalleled realism within a mix situation. The sounds of a range of different artificial reverberators can also be mimicked in this way, although any dynamically varying qualities of the original units cannot satisfactorily be emulated with convolution processing.

as well. If you see separate level controls for the reverb's initial echoes (Early Reflections or ER) and the remainder of the reverb decay (Reverb or Tail), then those can also be helpful, but they aren't essential.

There'll be a slew of other controls on a lot of plug-ins, but for typical small-studio operators—and indeed a lot of professionals—there simply aren't enough hours in the day to worry about what all of those do. (They do serve a useful social purpose, however, in distracting computer anoraks from ever getting any music finished.) As long as you choose a reverb plug-in that has a load of programmed presets, you can get along just fine without any other parameters, because the balance processing you already know about from Part 3 provides ample scope for further molding the sound of each reverb if necessary. So, without further ado, let's see how you go about designing and applying the first (and probably most important) type of reverb patch: blending reverb.

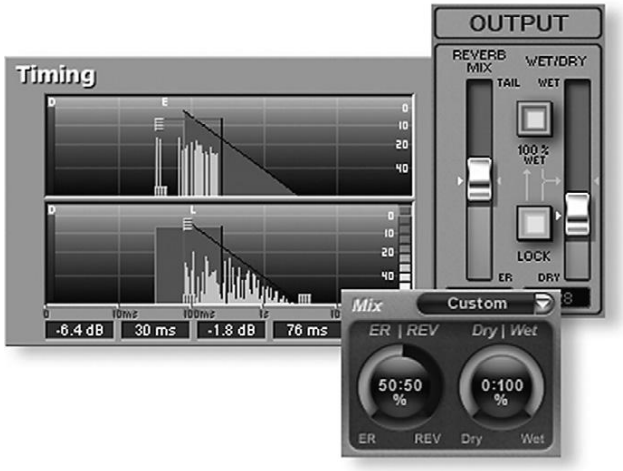

FIGURE 16.3
A number of reverb plug-ins provide the option to control a reverb's early reflections separately from its decay tail. Although it's nice to have access to such controls, you can actually get by fine without them most of the time.

16.3 REVERB FOR BLEND

The reason almost all small-studio productions need blend reverb is that a lot of their tracks are recorded as separate overdubs. Overdubbing avoids spill between different instruments in the arrangement, and though this causes fewer phase-cancellation complications, it also means that you don't benefit from the natural blending effects of sound leakage between different microphones. In addition, most modern recording sessions involve a lot of close-miking, and a close-mic usually picks up quite low levels of natural room reflections in most situations—and whatever reflections it does pick up will change a good deal as the mic is repositioned for recording different instruments. As a result, many tracks won't blend well enough with the mix, and sounds that should be in the background feel too close for comfort. Adding a common reverb to the unduly upfront tracks helps blend them together, pulling them further away from the listener and making them sound more as if they were recorded in the same place at the same time.

Play back the mix section you balanced for Part 3, and listen for any tracks that are having trouble blending—either they don't seem to connect with the mix as a whole, or they seem too close to you. Now select one of those tracks that has a comparatively full frequency range and preferably some transients too. (I often start with a mixed drum loop or a live drum kit's overheads, for example, but only if those parts don't seem to blend properly in their raw state.) Send at a decent level from this track to the reverb, and then solo the reverb return so that you only hear the reverb's added echoes, without any dry signal. Why kill the dry signal? Because that lets you really concentrate on the nature of the effect. Phil Ramone says, "One of the hardest things to teach somebody is

to listen to the device itself. Take out the source and listen to the device. You'd be amazed how crummy some of these things sound! They flange, they phase, there's cancellation all over the place."[3]

Picking a Preset

You're now ready to start hunting for a promising-sounding preset that feels like it fits the sound you want for the mix. This is probably the most critical part of creating a reverb patch, but because it's subjective and genre specific it's also difficult to advise on. The main thing to realize is that the part of a reverb that is primarily responsible for its blending effect is roughly its first half-second. So whenever you try a new preset, reduce its length straight away to home in on the blending characteristics. If you have independent level controls for Early Reflections and Reverb Tail, pulling down the latter should help too. It doesn't matter exactly how short you make the reverb for the moment—just shorten it enough to make it into a brief, well-defined burst of reflections rather than an echoey decay tail.

Take preset names with a large pinch of salt. Just because something says "Epic Snare Boosh!" doesn't mean that it won't serve as a good general-purpose blending reverb in your situation.

Beyond that, here are a few other general-purpose tips that may assist you in arriving at a good choice of preset:

- Take any preset names with a large pinch of salt. Just because something says "Epic Snare Boosh!" doesn't mean that it won't serve as a good general-purpose blending reverb in your situation. Nevertheless, do keep an eye out for any preset with "ambience," "early reflections," "short," or "dry" in its title, as there's a greater likelihood that it will be what you're looking for.
- Natural-sounding presets tend to do a better job for blending purposes than obviously synthetic reverbs (in other words, ones that don't really bear much relation to the sound of a real acoustic space). For example, any preset name containing references to plates or springs probably won't be that suitable.
- It can sometimes help to close your eyes and visualize the kind of space you want the blended sounds to inhabit. Picturing a real environment can help focus the mind here, although this may not help much if you're after a more otherworldly sound.
- Don't worry too much if the frequency-balance of the reverb doesn't seem right, because you can do a lot about that with equalization. What's most important is that the overall acoustic signature feels right. That said, patches where either of the frequency extremes are very pronounced are unlikely to be useful for blending purposes.
- If possible, steer clear of plug-ins or presets that produce a kind of metallic sound with clearly audible pitched resonances, especially in response to transients.

- Try to find a patch that is evenly spread across the stereo image and doesn't swerve drunkenly off to one side or the other as it decays. You'll be better able to hear the reverb decay if you mute and unmute the track send a few times.
- Don't rush! It's not uncommon to trawl through more than a dozen presets and change plug-ins a couple of times before you find a reverb sound that really seems in keeping with the production you're working on.
- Rely heavily on your nearfield monitors here, because they'll give the best overall tone and quality judgments, but be sure to check the effect in mono as well to ensure that you don't get any nasty tonal changes. Mono compatibility is one of the areas in which expensive reverbs tend to outclass cheap ones, so be especially careful here if you're working with budget or bundled plug-ins. "You really should check in mono," stresses Phil Ramone, "to make sure that what you're hearing is what you get. That's why a lot of television shows get so screwed up. They come in with these great effects, but when they check in mono… half the reverb goes away."[4]

Once you've found a likely contender, unsolo and mute the return channel to remind yourself for a few seconds how the production sounds without the reverb. Then unmute the reverb in the mix, fade it up so you can hear it clearly, and confirm that it's really worth pursuing. If you think it is, then the first thing to do is tweak the reverb's length so that it's short enough to tuck behind the dry sound in the mix without creating any audible reverb tail, but also long enough that it can deliver sufficient blend as you fade it up. In doing this, you're minimizing the degree to which the reverb affects the size and sustain of tracks it's applied to, which gives you more power to adjust blend independently.

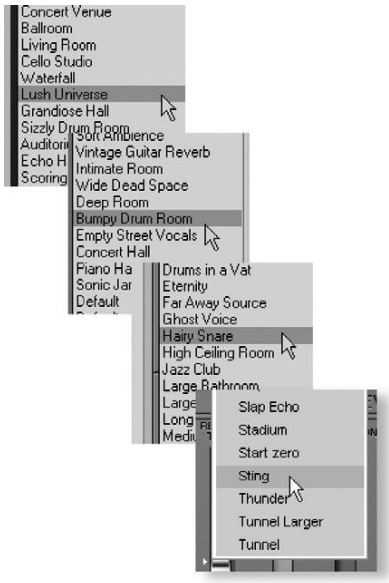

FIGURE 16.4
Just because a reverb preset's name is a bit silly, that doesn't mean it won't be useful. The only way to tell is to have a listen.

Tone and Spread Adjustments

Next you want to minimize undesirable tonal changes caused by adding the reverb to the dry signal, and this is probably easiest to do if you isolate the dry track and reverb return together. The tools and techniques for doing this are exactly the same as you used to adjust the combined tone of multimiked instrument recordings in Sections 8.3 and 11.3: timing and phase shifts, polarity inversion, and EQ. Clearly, though, you want to process only the effect signal in this case, otherwise you'll upset the balance you've already achieved for your dry track. Although some reverb plug-ins may include tools for adjusting any of these aspects of their output, it's usually just as easy to insert a favorite delay, phase-adjustment, or EQ plug-in from your collection into the reverb's return channel. Mostly these processing choices should need very little further explanation—you just have to listen for any unwanted tonal colorations introduced when you fade up the reverb return, and then rebalance the effect's frequency response to counteract them. However, the idea of delaying the effect signal, usually referred

FIGURE 16.5
A lot of reverb plug-ins have balance-processing features built into them, but it's usually easier and more flexible to use the ones you already know. All you have to do is insert your normal mixing plug-ins into the reverb return, as shown here in Cubase.

to as adding "predelay," does warrant a little extra discussion. "I think the predelay is the most important part of any reverb," says Alan Parsons, "and it's the first button I reach for."[5]

The thing about predelay is that it doesn't just adjust phase-cancellation effects between the wet and dry sounds; it also influences how far back into the mix the blend reverb appears to pull the treated track, subjectively speaking. The less predelay your blend reverb uses, the further this reverb can push sounds away from the listener. If you use no predelay at all, then the implication is that the treated instrument is stapled to the rear wall of the reverb's virtual room, whereas adding predelay rescues is from that uncomfortable predicament and moves it closer to the listener. My advice here is to start off with about 10 to 20ms of predelay, aiming on the shorter side for a more intimate-sounding space and on the longer side for a more spacious virtual acoustic, and then to refine the setting up to a few milliseconds either side of that initial setting for tonal reasons. What this means is that you can apply the blend reverb across all your tracks to some extent if necessary, without distancing the mix as a whole unduly from the listener. However, if you later discover that you're unable to blend or distance any track far enough, you may wish to set up a second blend reverb with less predelay specifically to treat that—this eventuality isn't common, in my experience, but it's useful to recognize it as a possibility. One other issue to bear in mind with predelay is that it reduces masking of the wet signal by the dry signal, so you can typically use lower reverb levels when you've dialed in some predelay, and this can improve the clarity of your mix.

The final characteristic of the reverb you may wish to tweak a little is the stereo spread. Having a wide stereo spread on a blending reverb is usually a desirable side

effect, because it helps the tracks in your mix feel as though they belong together by encompassing them in their acoustic environment, but there are some situations where you may wish to narrow the spread. For example, you might wish to blend drum spot mics with stereo room mics that either don't have a wide stereo picture or have been deliberately narrowed by your choice of panning—a wide-sounding blend reverb can seem rather incongruous in this case.

Balancing Blend Reverb

Once you've designed your basic blend patch, you need to balance it with everything else that's in the mix. You can't do this with solo buttons down, though, so deactivate those in order that you can listen to the reverb in its proper context. For the moment, fade down the single send you've added so far, and if you've been messing with the return channel's fader, reset it to unity gain. Without the reverb in the mix, reacquaint yourself with the blend problem for a few seconds to regain a bit of perspective, and then slowly raise the send level control as far as necessary to connect the dry track more strongly with the mix as a whole—or, to look at it another way, to distance the instrument far enough from the listener. The more reverb you add, the more the instrument will blend, and the further into the background it will move. As Dave Pensado notes, "If you want the singer to sound like she's standing behind the snare drum, leave the snare drum dry and wet down the singer."[6]

WHAT IF THERE'S TOO MUCH BLEND?

Lack of blend isn't the only blend problem that you might be presented with: you may also encounter tracks that blend too well. In other words, they've been recorded with too much spill/reverb so you can't actually bring them far enough forward in the mix. This can be almost impossible to fix, because it's tricky to get rid of recorded-in reverb. There are nevertheless a few processing tricks that can yield some success in "unblending," especially if used in combination:

- Brighten the instrument's tone with equalization or distortion processing, because we humans are psychologically programmed to interpret brightness as a clue that a sound source is close at hand. (High frequencies are so easily absorbed that they never survive a long-distance trip to the listener nearly as well as the midrange and low end.)
- Narrow the stereo image of stereo files, or even collapse them to mono, as the reverb will often phase-cancel more than the instrument itself.
- Increase the instrument's dynamic range in some way to dip the low-level background reverb tails between notes and at the ends of phrases. Simple expansion or limited-range gating may work to some extent, whereas transient processors can sometimes deliver the goods on percussive tracks if they offer the option of negative "sustain" gain. Multiband expansion or even dedicated noise-reduction may offer further improvements, but it'll usually be fairly slim pickings.
- You might make a last-ditch attempt at bringing the sound forward by layering in a dry MIDI instrument or sample alongside the recorded part.
- Kidnap the client's ornamental fern, and demand that the part be rerecorded as a ransom. For every day's delay, send the client a frond in the post.

Once you've got the most satisfactory setting you can, make a point of shifting your listening perspective to check that the reverb isn't masking any other instruments, just the same as you did when checking the masking impact of each newly added instrument during Part 3's balancing process. If some important frequency range is being obscured by the added reverb, then toss an EQ plug-in onto the reverb return channel to carve away at it. (You may then need to fade up the remaining reverb frequencies to regain the former degree of blend.) Pay particular attention to the region below about 300Hz or so, as this usually benefits from some degree of cut in most styles to preserve headroom and clarity for bass instruments at the low end and also to avoid muddy-sounding energy buildup in the lower midrange frequencies—the latter being one of the most common problems in small-studio mixes. "Often I'll EQ the reverb to attenuate lows," says Mick Guzauski, "so that… it doesn't cloud any instruments. Usually I'll roll off some low mids, around 200 to 300Hz."[7] The amount of EQ thinning that's suitable here, however, is quite genre-specific, so be prepared to revisit this setting after you've referenced your first draft mix against some commercial competitors.

Equalizing the reverb return also has other purposes, particularly if you'd prefer your reverb effect not to draw attention to itself—for example, in hard-edged or acoustic styles where audible artificial treatments might undermine the music's sense of rawness and emotional integrity. What most commonly gives the game away that you're using reverb is excessive high frequencies in the reverb return, because most musically useful natural acoustics rarely have lots of bright-sounding reflections. "If I'm going to use a reverb," says Elliot Scheiner, "I want it to be inconspicuous, so I don't use bright reverbs any more; my reverbs are generally darker."[8] A low-pass filter or high-shelving cut in the reverb return is, for this reason, very common when using any reverb, but it's almost a matter of course when the reverb's specifically targeted at

FIGURE 16.6
It's common to reduce the amount of low-end content in your reverb return (as shown in SSL's X-Verb here) to keep the low frequencies of your mix clear sounding.

providing blend. Exactly how you set the filter must, of course, be judged by ear in order to find the frequency contouring that best reduces the reverb's audibility while at the same time allowing it to fulfill its blending role. If you can't hear the artificial reverb as an effect in its own right, but the mix seems to thin out and fall apart when you mute the reverb return, then you're probably heading in the right direction. Ed Seay says, "You don't even have to hear it, but you can sense it when it goes away—it's just not quite as friendly sounding, not quite as warm."[9]

Sometimes equalization doesn't go far enough and dynamics also need to be employed. The most common circumstance that calls for this adjustment is where transients are causing the reverb to produce distracting flams or stereo ricochets, in which case it's worth trying out a dedicated transient processor in the return channel, placed before the reverb plug-in, to see if you can tone those down. Threshold-based transient processors are not usually effective in this role if more than one track is feeding the reverb, so threshold-independent designs are a much better choice. Sometimes frequency-selective dynamics prove necessary too, most commonly where vocal sibilance hits the reverb and sprays unnaturally across the stereo picture. A de-esser inserted into the return channel before the reverb is the obvious solution here, but a potential drawback is that threshold-independent de-essing isn't easy to find. If this causes a problem, your best option might be to send vocal signals to the reverb via an extravagantly de-essed intermediary channel. (Just don't send this channel direct to the mix or it'll sound like the singers have swallowed their dentures.)

Once you've applied your blend effect across all the channels that seem to require it, make sure to confirm the balances on your different monitoring systems at different playback volumes. It's also not a bad little reality check to bypass the reverb return for 10 seconds, envision how you want the track to blend, and then pop the reverb back in. Oftentimes this simple process will highlight some small balance problems or a previously unnoticed tonal coloration that needs attending to with further tweaks of the return EQ. You can also improve your ability to hear internal effect balances within a mix if you listen once or twice with the main tracks (perhaps bass, drums, and lead vocals) bypassed in various combinations, so that the foreground parts don't distract your concentration from background details—Geoff Emerick[10] and Alan Parsons[11] have both independently mentioned using this technique for their own mixes, so you'll be in good company.

All that said, it could well be that applying blend reverb to the first track sits it into the mix just fine without the need for any additional balance processing on the return channel. However, as you continue to identify insufficiently blended instruments and apply the same effect to those as well, the situation might easily change, so it's quite normal to keep tweaking your return-channel balance processing right up until the moment you declare the mix finished.

BLENDING BY OTHER MEANS

Reverb isn't the only means to blend a production. Short delays will also work admirably (as explained in the next chapter) and many modulation effects will distance instruments by making them more diffuse sounding. Double-tracking can blend an instrument or voice better too, even if you keep the double-track at such a low level that it's not really audible in its own right. Adding background noise to your mix will often improve the overall sense of blend, and if you don't fancy sourcing your own noise recordings, then get hold of a media sound-effects library and look for what are usually called "room tone" or "room ambience" files. Room tone is the sound of nothing happening in a room—not the most interesting thing to listen to, but in a background role it can really help make all the tracks in your mix feel as if they belong together. Tape hiss and vinyl noise have similar blending effects, especially when they're in stereo.

16.4 REVERB FOR SIZE

The biggest difference between reverbs designed for blend and size is that where the former is best provided by the earlier reverb reflections, the latter is best created by focusing the effect sound on the remainder of the reverb tail. As such, you can leave the reverb length of each preset unchanged during auditioning and turn down any separate level control you have for early reflections. Presets where the reverb onset has a clear attack tend to work less well for this application than presets that start more softly and slowly. Beyond these suggestions, much of the previous advice about choosing a preset still holds:

- Beware of unnatural-sounding presets, as these will have trouble creating the sound of a larger space convincingly. CPU-light plug-ins will typically sound less natural than more computationally hungry algorithms.
- Feel free to ignore the preset names with impunity—the main goal is to try to imagine the space implied by each preset and decide whether it's the right kind of space for your mix to exist within.
- Don't be too concerned about tonal imbalances as long as there aren't nasty metallic resonances.
- Check that the stereo picture is fairly evenly spread, and assess the mono compatibility.
- Don't hurry the selection process, and make sure you ratify your choice properly with your different monitoring systems.
- When you've got a promising patch, mute it, recalibrate your ears to the mix as is, and then fade it up to confirm that it's actually what you're looking for.

Any reverb you create for size enhancement will inevitably also create a small amount of blend too (much as any blend reverb will also give hints of size), but you can reduce this "mission creep" to a useful degree by giving your size reverb a good dose of predelay—anything from 50ms upward. It's because of

this longer predelay that reverbs with an attack "bump" won't usually work in this role, because the hard reverb onset can potentially make transients sound like they're flamming. It's partly for this reason that some engineers opt to use a tempo-related predelay time, because then the reverb attack following any percussive hit will be masked by the following one, making the bump less audible. Glen Ballard provides another good reason for increasing the predelay for any longer reverb: "I'm always interested in hearing the vocalist's words… so I like to have the dry signal a little bit clear of the effect, just for the articulation. If it's a really slow song, say a ballad, then the pre-delay will probably be longer."[12]

Size reverb will also inevitably add sustain, and though this may in general be a good thing, it's wise to try to make the tone of that sustain fairly neutral if possible, because you may wish to apply this effect to many sounds in your arrangement if you're after any kind of "all in one room" feel to the production. Given the long delay time, phase cancellation between the wet and dry sounds should be minimal, so any tonal adjustment can be carried out simply with EQ—solo a full-range instrument along with the reverb return, and then fade the reverb up and down to highlight tonal changes in the instrument's sustain that you might want to address. You may find notch cuts to be helpful here if the reverb turns out to emphasize any dissonant pitched resonances. As with blend reverb, the stereo spread can normally be left as is.

Balancing Size Reverb

With a suitable reverb patch on the go, deactivate those solo buttons again so that you can get on with balancing the effect within the context of the full mix. The procedure is similar to that described for blend reverb, the main difference

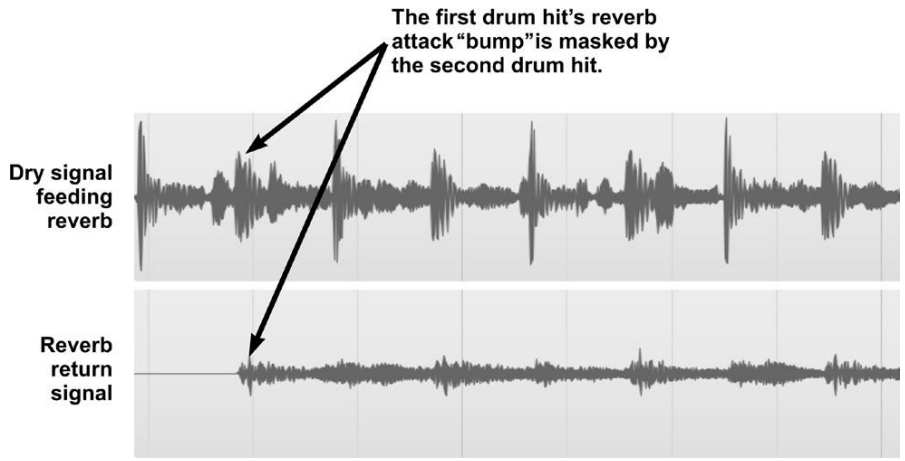

FIGURE 16.7
Reducing the audibility of a reverb's attack "bump" using a tempo-matched predelay time.

being that fading up the reverb here gives the impression that an instrument is within a bigger space than it was recorded in. The more you apply, the clearer the impression of that larger acoustic becomes. The more you fade up the reverb for any specific instrument, the more audible the reverb becomes and the greater the additional illusion of relative size and power. With the size enhancement will inevitably come a certain amount of extra distancing of the tracks that contain the effects, as well as some additional sustain, but these should be secondary effects that aren't as strong as the size illusions.

There is an additional subjective decision to make, though, with size reverb: Should you bring all the instruments within a single space, for a more realistic presentation, or be more creative and place individual instruments or groups of instruments into contrasting acoustic environments? The latter is far from natural or realistic, of course, but nonetheless provides excellent opportunities for musical contrast: your lead vocal might be dry as a bone so that it's right up at the front of the mix, the drums might be in what sounds like a large wood-paneled room, and the strings and backing vocals might float ethereally within a chapel-like environment—the sky's the limit, really. Creative choices like this aren't really within the scope of a general-purpose mixing primer, but I will offer one small piece of advice here that has served me well in the course of many different mixes: try to match the nature of each different space to the sonic character and artistic intent of the sounds it will be applied to. So don't put a smooth concert hall on aggressive punk drums if a ragged "in the garage" room sound fits their attitude better. Don't stick an angelic boys' choir in a long corridor if a spacious cathedral sound supports the serenity of the choir's musical lines more. Above all, don't put a bagpipe solo into a concert arena when it really belongs in an underground bunker.

The longer the reverb, the greater the potential for masking problems, so it's even more important with size reverb that you keep shifting your listening perspective to assess the impact of the reverb on the clarity of your dry tracks, especially at low frequencies. Within most mainstream productions, reverb tails on the low end of bass instruments are usually counterproductive—what extra impression of size you get from the seismic rumblings is more than offset by a loss of headroom and note definition. My recommendation is to high-pass filter the reverb return to keep the low frequencies tight and controlled, and also to be conservative with the amount of size reverb on your bass instruments in general. Another thing to be careful of is adding lots of this kind of reverb to sustained stereo chordal parts such as synth pads. The nature of such parts makes it difficult for any long reverb to enhance them much, so all the reverb tail actually ends up doing is blurring the chord changes, making it feel like your keyboard player can't keep time. (And after all that work you did quantizing him too.)

Whether you want your size reverb to be audible as an effect in its own right is another judgment call that depends on the purpose of that reverb in

SPECIAL EFFECT REVERBS

This chapter is primarily concerned with reverbs that serve more technical purposes, but in some productions a reverb might also be used as a feature in its own right, a special effect designed to make a statement. If you're using a reverb in this way, there's not a tremendous amount of useful advice I can offer, simply because everything depends on your personal idea of what kind of effect you'd like to hear. My main recommendation is to consider incorporating that particular effect at the balancing stage, treating it exactly like any of the normal recorded tracks in your mix project and balancing it with just the same care. Bear in mind, though, that the effect might not rank nearly as highly in terms of "sonic importance" as the dry track that is driving it, so it may only be appropriate to add it to the mix further down the rank order. In addition, you'd be well-advised to ask yourself tough questions about whether all the aspects of the reverb in question are absolutely essential to its artistic impact. If you don't, then your mix clarity may suffer at the hands of unnecessary blend, size, tone, sustain, or spread side effects.

Also bear in mind that the impact of any special effect diminishes with familiarity, so it may be counterproductive to leave it in the mix all the time. "I'd rather flick it on and off throughout the song in two or three places," recommends Al Stone. "That's what an effect should be. If it's on all the time it cancels itself out."[13] Guy Sigsworth shares similar views: "I'm very fond of these electronic pieces Stockhausen did in the 1950s.... One of the few effects he had to play with was reverb, and he did a lot of things where the first note's completely dry, the second one as reverbed as possible, the next one somewhere in the middle. The reverbs are very dynamic, they're not like just some vat you dip the whole thing in, they move around, they're very agile, they're like a nervous animal. I really like that, so I try to get that hyperactive approach to space."[14]

the production. Where a single reverb is operating over most of the tracks in the arrangement to enhance the apparent size of the virtual venue, then keeping it understated is usually sensible, in which case you'll want to follow similar return-channel EQ and dynamics policies as when using blend reverb. However, where multiple size reverbs are being used more creatively, then making the sound of some of these conjured-up spaces less bashful might help with the musical contrasts within the arrangement—just be doubly careful of undesirable masking consequences from the reverbs in this case.

Whichever tack you adopt, do still go through the usual objectivity test with each reverb once most of the sends to it are operational: bypass the return for a few seconds, recreate in your mind the size illusion you're after, and then reinstate the reverb to see how it compares to that ideal. If dropping out some of the more important instruments for a while helps you refine the subtler reverb balances, then give that a razz as well. One matter to devote particular

concentration to while finessing the overall sound of your size reverb is whether you've chosen the best combination of reverb length and overall reverb level. If the reverb is too long, then you may not be able to use enough of it to create a believable illusion of a larger space without it washing out the whole mix and obscuring all the nice mix details; if the reverb's too short, then you may not be able to enhance the apparent size of your production as effectively as you'd like to.

16.5 REVERB FOR TONE

Where it can frequently be appropriate to have sends to a blend or size reverb coming from most of the channels in your mix, reverb patches dedicated to tone or sustain enhancements will usually be applied only to small groups of tracks. This is because these last two reverb types have more in common with EQ than they do with reverb, and it should be clear from Part 3 of this book that an equalization setting that works for one instrument is unlikely to work well for another. Indeed, the similarity of such processing to EQ was the reason I introduced the idea of tonal/sustain reverb back in Chapter 12—it's much more sensible to do any subjective timbral shaping of a sound early on in the balancing process, rather than leaving it until the balance is complete, whereupon any tonal changes could easily upset the whole applecart. Now that we've discussed reverb in more detail, however, let's look further at how to get the best from these more coloristic enhancements.

FIGURE 16.8
Although quirky little CPU-light plug-ins like Voxengo OldSkoolVerb (*top*) or GSi TimeVerb (*bottom*) may not sound natural for enhancing blend or size, they can really add character when you're looking for tone or sustain alterations.

A reverb effect primarily for the purpose of altering the tone of a treated track is best shortened to avoid obvious reverb tail, much as it would be in a blend reverb, because this immediately minimizes size and sustain side effects. So when you're looking through your plug-in presets you'll want to keep the effect's length under control to get a decent impression of each reverb's tonal qualities. The similarity with blend reverb ends there, though, because the most usable kinds of presets will typically be those with an unnatural sound—the less they sound like real acoustic spaces, the less they'll blend the track as a side effect. Hideous frequency-response imbalances and horribly clangorous resonances can also be an advantage, because those characteristics will be more likely to bring about truly meaningful changes in the timbre of treated sounds. Any rubbishy old plug-in with a wafer-thin CPU load is also right back on the menu for similar reasons. Jack Douglas observes, "Sometimes the weirdest things—like old spring reverbs—can sound really phenomenal in the mix. By itself it's going to sound awful… but use it right—color it a little bit, filter it, EQ it—and it's going to sound cool in your mix."[15] Manny Marroquin shares his taste for springs: "They're cool-sounding. You can make [the reverb] short and tight with a gate and it adds tone and depth to the sound without washing it out…. When you have one thing in a crowded mix you really want to bring out, but adding EQ would make it sound harsh, put it through a spring reverb."[16]

If a tonal reverb is substantially wider than the image created by the dry signals, or in a different location in the stereo image, then the timbral change may be diluted and the treated sound will appear less upfront in the mix.

What you also need to realize is that the sound of a particular preset's echoes on its own is only part of what you're looking for, because a big proportion of the final tonal change available here will actually result from comb filtering between the wet and dry sounds. For this reason, you should always make sure that the dry sound is mixed in alongside the effect when auditioning presets for a tonal reverb patch, and bear in mind that the relative levels of the dry and wet signals will impact on the severity of the phase cancellation. Moreover, once a preset has been found that flatters the dry signal, you may be able to refine it further by applying predelay and phase/polarity adjustment to the return channel—note that predelays under about 10ms will usually give you more powerful comb-filtering effects. Equalization of the return channel also has a lot to offer. Manny Marroquin states, "I listen for the frequency where the reverb matches the input signal and tweak that."[17]

Ideally, the goal here is for the reverb not to be heard, so that it fuses with the dry sound in the listener's perception to provide the greatest tonal changes with the fewest blending side effects. One of the challenges to achieving this objective is finding the right combination of reverb level and reverb length—as the level of the reverb increases, its length will often have to be decreased if you want to maintain the same level of audibility. It's also wise to be careful of the stereo width and placement of your tonal reverb return, because if the effect

is substantially wider than the image created by the dry signals, or in a different location in the stereo image, then the timbral change may be diluted and the treated sound will appear less upfront in the mix. You're probably already firmly in the habit of switching monitoring systems by now to ratify tonal mix decisions like this, but do ensure you check mono compatibility in particular to avoid any nasty surprises—as I mentioned earlier, cheap and nasty reverbs can really fall on their asses in mono.

The danger with tonal reverb is that adding it will usually increase the overall level of the treated instrument, especially if you're using similar levels of wet and dry signal to exploit the full possibilities of comb filtering. For this reason, it's possible to mislead yourself if you simply switch off the effect in an attempt to judge its value. A better approach is to find a reasonable fader level for the processed sound in the balance and then mute the dry and wet sounds together for a few seconds. When you've had a chance to appreciate the state of the mix without those channels, try to bring to mind what you want the composite to sound like and only then reintroduce them to the mix. You're more likely to be dispassionate about the reverb's merits that way.

Once you're confident that you've arrived at a combination of dry and wet signals that truly improves the character of selected instruments in the mix, you're still faced with the task of balancing these rejuvenated sounds within

FIGURE 16.9
A typical tone reverb setting, running in Christian Knufinke's SIR2 convolution reverb plug-in: a characterful impulse response with a short decay time, a sub-10ms predelay, narrowed stereo image, and heavily sculpted frequency response.

the mix as a whole. Because the tonal reverb effectively becomes an integral part of the sound, it stops functioning as a traditional send-return reverb effect, and it makes much more sense to treat the dry and wet channels in the same way as you would the individual microphone channels in a multimiked instrument recording. So I normally end up routing them both to a communal group channel for global control, which allows me to process the complete sound for balance reasons without altering the fastidiously honed phase relationships between the constituent channels. As with multimiked instrument recordings, however, there may be further creative potential in processing the individual channels, but you can run into difficulties here if you add the tonal reverb to more than one track, because if you process the effect channel for the benefit of one of the dry tracks, it may not suit the others. It's not much extra effort to gain extra control here, though, by creating separate instances of the same effect for each treated track. You may also be able to reduce the audibility of the effect in this way, by replacing a stereo reverb patch with several mono instances, each panned to the location of the dry track that feeds it—a tactic that both Al Schmitt[18] and Elliot Scheiner[19] have employed successfully. Of course, this does inevitably multiply your plug-in count, so this is an additional reason to favor algorithms with a small CPU appetite in tonal roles.

16.6 REVERB FOR SUSTAIN

What tonal reverb is to blend reverb, sustain reverb is to size reverb—in other words, a longer reverb designed for creative timbral control rather than the simulation of natural acoustic reflections. As such, setting up a sustain reverb involves many of the same steps already discussed with relation to tonal reverb in Section 16.5:

- You'll usually only apply dedicated sustain reverb to small groups of instruments, and it's best to do this at the balancing stage if possible.
- You can only choose presets effectively while listening to the wet signal alongside the dry.
- Reverbs that sound too natural will tend to have more unwanted side effects, in this case primarily adding size enhancement, so don't be afraid to wheel out plug-ins that would curdle milk at 20 paces under normal circumstances.
- Your aim is to get the dry and wet sounds to meld into one perceptually. If they don't, the effect will be more likely to muddle the mix, increasing size-enhancement side effects and masking the details of other instrument sounds. Within this context, the reverb's level, length, and stereo spread/position all require due consideration.
- Resist evaluating the success of sustain reverb by simply muting/unmuting it, otherwise you'll fall foul of loudness bias. Instead, drop the whole sound out of the mix, give yourself a few seconds to reacclimatize and rebuild your mental image of how the treated sound should appear in the mix, and then switch the instrument (and its effect) back in.

- Once you've achieved the best sustain enhancement you can manage, you can effectively pretend that the dry track and wet return are the individual mics of a multimiked instrument recording for balancing purposes.

GATED REVERB

It's a shame that the excesses of the 1980s have given gated reverbs such a dreadful reputation, because many engineers now instinctively shy away from them even when they're useful. A couple of specific situations come to mind: when you want to achieve the maximum sense of size in a production, but without ridiculously long reverb tails obscuring all the mix details or blurring chord changes; and when you want to add a short, dense tonal or sustain reverb that doesn't gradually decay, an effect that's particularly good for bulking out thin-sounding drums/percussion. Although traditional gated reverb was generated using a gate acting on a normal reverb, in my experience the later "fake" gated-reverb effects (constituting a short burst of echoes that don't decay in level) are much more usable for mixing purposes because their additions are threshold independent—a quiet drum hit will receive the same length and density of reverb burst as a loud one.

There are a few further considerations with sustain reverb, though. For a start, it's a smart idea to use a good dose of predelay (I'd suggest 25ms or more) to remove phase cancellation from the equation. That way, you'll minimize direct EQ-style tonal changes, which gives you more independent control over the sustain enhancement as well as making any balance processing of the reverb return respond more predictably. However, as with size reverb, that increase in predelay brings with it the possibility that a clearly defined reverb onset may draw unwarranted attention to itself in the mix, so you either have to avoid bumpy presets or use a tempo-related predelay time.

16.7 REVERB FOR SPREAD

The stereo enhancement aspect of reverb is difficult to separate from its other functions, so most of the time it's best to think of it more as a bonus free gift rather than the main purpose of the effect. There are situations, however, where the stereo aspect of the reverb may be the main appeal. Here are a few examples:

- A classic effect for vocals involves a short fizzy reverb that adds density and stereo width to the upper frequency range. The lack of tail minimizes size and sustain contributions, whereas the combination of predelay and high-frequency bias avoids any sense of blend and restricts the tonal impact to a generic high-frequency lift.
- Any effect with a reverb tail can be used to "paint to the edges" of the stereo field in a generic way. Reducing early reflections and adding substantial predelay avoids adjustments to blend or tone; a smooth, bland, but

nonetheless unnatural preset choice reduces size-related side effects; and balance processing of the return channel minimizes any undue sustain emphasis of individual frequency ranges and reduces the overall audibility of the reverb as an effect in its own right.

- Short reverbs can operate in a similar way if you use characterless but synthetic-sounding presets and increase predelay to scale down blend and tonal enhancements. Be careful of flams, though, if your predelay goes beyond about 20ms.

16.8 JUGGLING REVERB ENHANCEMENTS

Learning how to create and apply each of the different specialized reverbs I've discussed here serves two purposes. In the first instance, it means that you can use a combination of these effects to add exactly the right amount and quality of each enhancement without drowning in unwanted side effects. Whenever small-studio mix engineers tell me that they're having trouble finding the right reverb levels in their mix, it's invariably because their choice of plug-in isn't providing a suitable combination of enhancements. By the time there's enough blend, the tone warps; when you've managed to add some flattering sustain to your piano, it sounds like it's in a railway tunnel; and in the process of adding a sense of space, you blend things too much, so everything sounds a million miles away. Dedicating specialized reverbs to specialized tasks combats this all-too-common frustration.

Whenever small-studio mix engineers tell me that they're having trouble finding the right reverb levels in their mix, it's invariably because their choice of plug-in isn't providing a suitable combination of enhancements.

UNMASKING WITH REVERB

Although reverb can often cause additional masking problems in a mix, it's worth realizing that it can sometimes actually counteract the effects of masking. One way it can achieve this outcome is by increasing a masked instrument's sustain such that the wet signal trails on in time beyond the end of the specific musical event that is masking the dry signal. The stereo width of a reverb patch can also unmask instruments by virtue of the wet signal emerging from the "shadow" of an instrument that is masking the dry signal. For instance, if a central shaker in your arrangement were masked behind simultaneous central kick drum hits, you could use reverb to unmask it, either by extending the shaker's sustain past that of the kick or by widening the shaker sound on either side of it in the stereo field. In addition, tonal reverb might also unmask an instrument by adjusting its frequency balance, in much the same way equalization could.

The second advantage of knowing which reverb characteristics affect the prominence of which enhancements is that it enables you to adopt a less regimented approach if the practical limitations of your studio demand it. Working in a hardware setup, for example, might limit you to only a small number of reverb processors, in which case it may be necessary to use each individual effect to cover more bases in terms of the possible enhancements. Once you understand how to minimize and maximize each individual reverb enhancement, it isn't a huge step to create reverb patches with desirable proportions of several enhancements simultaneously. If you have personal aesthetic reasons for using an idiosyncratic selection of effects, perhaps if you're after a specific retro-tinged sound, then an awareness of whether/how a specific processor's enhancements can be controlled is also extremely useful in determining how to get the maximum value out of its unique contribution to your production. The same applies if there are any recorded reverbs in the mix, whether in the form of ambient microphone recordings or "print it now because we'll never find it again!" effects bounces.

To finish up this discussion of mix reverb, my final piece of advice is this: don't use reverb as some kind of involuntary reflex. To quote Steve Albini, "[Reverb] is a thing that's done *pro forma* a lot of the time. [Engineers] put it on because they feel they're supposed to. I've never had that response. I'll wait until someone says, 'That sounds weird,' and then I'll try reverb."[20] Numerous commercial records have precious little reverb on them at all, because they already have sufficient blend, size, sustain, tone, and spread as it is. Many other modern productions apply reverb only selectively because a lack of some of reverb's enhancements can support the overall intent of the music—for example, a lot of music in pop, urban, and electronica styles benefits from some elements being poorly blended, so that they're right in your face at the front of the mix.

CUT TO THE CHASE

- Reverb enhances several aspects of a treated sound at once, the most important for mixing purposes being blend, size, tone, sustain, and spread. To use reverb effectively at mixdown, you need to gain some independent control over each of these enhancements so that you can achieve exactly the right combination.
- To create a blend reverb, shorten a natural-sounding stereo preset to remove its decay tail, and add around 10 to 20ms of predelay. Address unwanted tonal colorations by adjusting the predelay, phase, polarity, and equalization of the effect.
- To create a size reverb, choose a natural-sounding stereo preset preferably with a soft reverb onset, and add around 50ms or more of predelay. Address unwanted sustain colorations by equalizing the effect.
- To create a tonal reverb, shorten an unnatural-sounding preset to remove its decay tail and refine the sound with predelay, phase/polarity, and EQ adjustments. Pay careful attention to the length and overall level of the

effect, and consider whether its stereo width and positioning suit that of the dry signal.

- To create a sustain reverb, choose an unnatural-sounding preset, preferably with a soft reverb onset, add at least 25ms of predelay, and refine its sound with EQ adjustments. Pay careful attention to the length and overall level of the effect, and consider whether its stereo width and positioning suit that of the dry signal.
- To create a spread reverb, try an unnatural preset with a bland, uncolored tonality, and do your best to use predelay and balance processing to minimize blend/size/tone/sustain side effects.
- Blend and size reverbs will often be added to many different tracks in an arrangement and are usually best applied after the complete balance is in place, whereas tonal and sustain reverbs (and indeed more ostentatious reverb special effects) tend to be more useful for very small numbers of tracks and are often better applied during the initial balancing process itself.
- Blend and size reverbs can be balanced in the mix almost as if their returns were independent instruments. Equalization can reduce unwanted masking, avoid low-end clutter, and reduce the apparent audibility of the effect, whereas de-essing and transient processing applied to the reverb's input signal can also help make the effect more understated. Tonal and sustain reverbs are usually better treated as an integral part of the dry tracks they are treating—almost as if the dry and wet signals were individual mic signals within a multimiked recording.
- Despite their checkered history, gated reverbs can still be useful, especially the "fake" designs that create an unnatural burst of dense, consistent echoes.
- Although reverb has the potential to mask elements of a mix, it can also help unmask instruments in certain situations by lengthening or spreading the dry sound or by altering its tone.

Assignment

- Do a survey of all the reverb plug-ins on your DAW system, and ensure that you have at least one fairly realistic reverb as well as a couple of character devices. If you feel restricted here, check out third-party alternatives.
- Return to your balanced mix section from Part 3 and see if you can get extra sonic mileage out of any of the tracks using tonal or sustain reverb patches. Try at least a couple of each to get an idea of the possibilities. If they produce something worthwhile, you may need to rebalance to accommodate the additions.
- Set up at least one blend reverb and one size reverb in order to experiment with how these enhancements affect the character of your mix. Don't forget to spend time balancing the return channels, with processing if necessary.

www.cambridge-mt.com/ms-ch16.htm

This page intentionally left blank

CHAPTER 17

Mixing with Delays

A delay effect creates patterns of echoes that are typically much simpler than those of a reverb. Although most small-studio mix engineers seem to favor reverbs for mixing purposes, delays can actually be more useful because of the way they deliver reverb-style enhancements in more precisely targeted ways—think of delay as a sniper rifle compared with reverb's shotgun! For many upfront modern productions, in fact, delays may be more suitable than reverbs, simply because they take up less space in the mix and can be made less audible in their own right. For example, most of what people think is reverb on Spike Stent's productions is actually delay.[1]

17.1 ESSENTIAL DELAY CONTROLS AND SETUP

When faced with a delay plug-in, the two main controls to worry about are Delay Time, which sets the time-delay between the dry sound and the first echo, and Feedback Level (sometimes called Repeats or Regeneration), which determines how many subsequent echoes follow the first and how quickly they decay over time. If you're comfortable with these two controls, then you can get good use out of plug-in presets and pretty much ignore other more detailed parameters unless you feel particularly inclined to pamper your inner nerd. Delay and reverb are both echo-based effects, so it should be little surprise that using each of them at mixdown is similar in many respects:

- A send-return effect configuration will give you more control over the sound of the delay effect that an insert setup.
- You can use delay effects to enhance blend, size, tone, sustain, and spread, although the effects in each case typically take up less space in the mix than those created by reverb.
- Simple, clean echoes tend to work best for blend and size enhancements, whereas more characterful algorithms (such as digital emulations of analog echo devices) deliver more varied tone and sustain changes.

Mixing Secrets for the Small Studio.

FIGURE 17.1
Emulations of old-fashioned tape-based delay devices are ten-a-penny, and their inherently unnatural (though appealing) sonic character makes them well-suited for delivering tonal and sustain enhancements.

- Clearly audible creative delays are best introduced while balancing, rather than trying to fit them in once the balance is complete.
- Don't forget to pay sufficient attention to balancing your delay returns, both in terms of level and in terms of equalization/dynamics processing.

If you're going to create a delay for blend purposes, then a single 50 to 100ms short "slapback" delay with very little feedback is a good choice. This is a long enough delay time to avoid the tonal side effects of comb filtering, but it's still usually short enough to gel with the dry sound fairly well, provided that the effect levels are quite low and balance processing is used to camouflage the wet signal as much as possible. Don Smith is just one fan of this kind of patch: "It adds a natural slap like in a room, so to speak, that maybe you won't hear, but you feel."[2] For size purposes, a longer delay with some feedback makes more sense, but otherwise the role of balance processing remains much the same.

We dealt with the idea of using delays for tone and sustain back in Chapter 12, because their use is often better dealt with at the balancing stage. However, now that we know more about echo-based effects in general, let's be a bit more specific. A tonal delay will produce the most dramatic effects when its delay time is well within the comb-filtering zone (under 20ms), although unnatural-sounding analog style delays will also inevitably color the tone even beyond that range. You may be able to get away with adding a little feedback if you want to emphasize the resonant "ringing" of the delay patch, but if so then you should consider whether you need to match the pitch of that resonance to the track you're processing—you

can tune it by finessing the delay time. Delays for sustain purposes are more likely to use substantial feedback, as well as longer delay times.

The processing on a blend or size delay's return channel will usually focus on balance issues, much as with reverb patches fulfilling similar functions, although many engineers also send from their delay returns to reverb effects in order to push the delay repeats back behind the dry signals in the mix, making them sound more like natural echoes. It's also common for delay effects to be ducked in response to lead vocals/instruments (as I mentioned in Chapter 14) so that the delay repeats don't interfere with note definition and lyric intelligibility. With delays for tone and sustain, there's much to be gained by processing the effect returns quite severely with EQ or the other tonal effects discussed in Chapter 12—remember that it's not what they sound like on their own that matters as long as they enhance the dry track. Making such effects sonically quite different to the dry tracks has an additional advantage that it reduces blend and size enhancement side effects, which can be useful if you want your production to feel tight, compact, and right in the listener's face—a favorite trick of Spike Stent's, which you can hear all over his mixes for Massive Attack.[3]

Making delays sonically different to the dry tracks has the advantage that it reduces blend and size enhancement side effects, which can be useful if you want your production to feel tight, compact, and right in the listener's face.

With size or sustain delays, the issue of how their longer delay times lock in with the song's tempo bears a little further thought, because matching a delay to tempo effectively makes it less audible—the echo on any given rhythmic event will be masked to some extent by the next rhythmic event. A tempo-synced delay can therefore be faded up much higher in the mix before it begins to become obviously noticeable, which frequently makes this kind of patch preferable for workhorse size/sustain-enhancement duties at mixdown. Most DAW delay effects now include tempo-matching functions as standard, allowing you to specify delay times in terms of musical note lengths, but if not then you can still work out tempo-related delay times in milliseconds very easily: divide 60,000 by the beats per minute (BPM) tempo value of your production, and then try different multiples and fractions of that value. So a tempo of 120bpm will give you a value of 500ms, for instance, and times of 1000ms, 250ms, or 125ms may also be useful—as indeed may triplet values of 167ms and 83ms.

It's not uncommon to have several tempo-synced delays at mixdown, as each delay-time note value may suit some dry tracks more than others. You may also find that a slight lengthening of the delay time beyond the strict tempo-matched value gives a more musical result in some cases. "I either set delays right on the beat, or… a little slow," says Neal Avron. "I don't like delays that rush. In more vibey songs, a slightly behind-the-beat delay can make things sound a little bit more laid-back."[4] Marginally lazy echoes can also be necessary where a performer is wandering a little ahead of the beat from time to time.

FIGURE 17.2
Plug-in delay effects like Steinberg's Cubase ModMachine and the Interruptor's Bionic Delay make it easy to synchronize your delay times to your project's master tempo. This is useful if you want the delay effects to sink into the mix and be less audible in their own right.

Delay effects that aren't tempo-synced can be much lower in the mix while still remaining clearly audible, so they're better for more creative applications where they enable the listener to appreciate the timing complexities without the effect trampling over everything else in the arrangement. The practical problem with unmatched delays, though, is that they quickly start to undermine the groove of your track if you use them too much, so if you're using delays as a statement within rhythmic styles, then it's usually better to take a kind of fence-sitting hybrid approach by choosing polyrhythmic delay times. These are tempo related but don't correspond to simple note lengths. Common examples are three-eighth-note and three-sixteenth-note delays. Polyrhythmic delays are the stock in trade of many dance styles, because they allow you to blow the average raver's chemically addled mind with lots of trippy-sounding audible delay repeats while safeguarding the essential rhythmic clarity of the groove.

17.2 USING DELAYS IN STEREO

Another reason why delays are often more useful than reverbs for modern styles is that they don't smudge the stereo field in the way a reverb does. As a result, delays typically leave a mix more open sounding and are preferred by many engineers.

There are different ways you might want to position echoes in the stereo field for different purposes. Blend and size delays tend to be most effective when you use a true stereo setup, in other words a delay effect where the stereo positioning/spread of echoes in the delay return matches that of the mono and stereo dry signals that are feeding the effect. In this scenario the delay repeats reinforce the stereo positioning of the mix and therefore draw less attention to themselves, which is ideal for subtle enhancement applications. In my experience, implementing a true-stereo setup in some DAWs is about as entertaining as brushing your teeth with an angle grinder, but try to persevere with it, because it's definitely worth the effort. (The stereo delay setup, that is, not the angle grinding.) Mono-input, stereo-output delay effects that add a sense of stereo width tend to draw more attention to themselves as artificial effects and may also confuse the imaging of stereo tracks, whereas simple mono-in, mono-out delays are less successful as global send effects because the mono return effectively narrows the mix and clutters the center of the stereo image. Mono delays do have their uses, but it's with tonal and sustain processing of individual instruments where they really come into their own, especially where the delay return is panned to match the dry signal being treated.

MULTITAP DELAY

Multitap delay effects occupy the middle ground between simple delay effects and reverb. What they allow you to do is independently specify the timing, level, and (usually) stereo placement of a handful of separate echoes. On the one hand, you can use these facilities to implement normal delay effects, but on the other hand, you can simulate something akin to a reverb's early reflections pattern. Where any specific multitap delay patch resides on this continuum will dictate how it's best dealt with from a mixdown perspective.

Of course, delay effects can also enhance the stereo spread, just as reverb can, if you can get your hands on a delay plug-in that gives you control over the positioning of the different echoes it generates. There are a couple of things to be aware of when doing this, though. First, panning short delays can disguise comb-filtering effects when listening in stereo, so you should confirm the mono compatibility of any such patch carefully. Second, there's a danger that panning individual echoes will skew the overall stereo image of the delay effect to one side, which might not be desirable. One response to this dilemma is to have alternate delay repeats opposition-panned so that they balance the stereo effect to an extent—a setup called "ping-pong" delay, which can be extremely effective for polyrhythmic patches in particular. Another tactic is to use a single echo panned to the opposite side of the stereo panorama than the dry signal. Whatever you try, though, you should always ask yourself whether the stereo character of your delays is distracting too much attention from more important elements of the production as a whole.

CUT TO THE CHASE

- Delay effects offer similar mixdown enhancements as reverb, while taking up less space in the mix. Although the echo patterns they create are usually simpler, you can treat them in much the same way as reverb from the perspective of setup and balance processing.
- Delay times that are matched to tempo will tend to sink into the mix, making them ideal for mixdown enhancements. Unmatched delay times will pop out of the mix at lower levels, so they are better suited to more ostentatious creative applications. However, if unmatched delay times are undermining the rhythmic flow of your production, then try out polyrhythmic delays instead.
- A true-stereo effect configuration will work best for blend and size delays, even though such a setup can be difficult to create in some DAW systems. Mono-in, stereo-out and simple mono delay effects tend to suit tonal and sustain roles better, but they may also be applied for the sake of deliberate stereo enhancement. If so, keep a careful ear out for mono compatibility and stereo balance.

Assignment

- Do a survey of all the delay plug-ins on your DAW system, and check that you have at least one fairly clean digital algorithm as well as a couple of characterful analog-modeled devices. If you feel restricted here, check out third-party alternatives.
- Return to your mix, save it, and open a copy. Mute any reverb effects you've used, and experiment with delay patches instead to see if they can achieve similar enhancements with fewer unwanted side effects. Once you've seen what different delay patches might be able to offer in each of the roles, save your mix again, open another copy of it, and remove all the effects sends, leaving the delay and reverb returns in place. Now try to achieve the best balance of delay and reverb effects for each channel. Once you've achieved this, save the mix once more, and then compare it with both of the previous versions to gain a full appreciation of what delays and reverbs can do at mixdown.

http://www.cambridge-mt.com/ms-ch17.htm

CHAPTER 18

Stereo Enhancements

During the discussion of panning in Section 8.2, I argued that stereo width is far more important to many mainstream productions than exact stereo positioning, a view that is borne out by the enormous variety of tricks audio engineers have come up with to adjust image spread. This chapter details a selection of these techniques in order to clarify each method's benefits and drawbacks.

18.1 ARRANGEMENT TWEAKS

Some of the most effective stereo enhancements at mixdown can actually be achieved surreptitiously without loading a single plug-in. My most common trick for inflating the sonics of Mix Rescue submissions, for example, is to generate fake double-tracks, which can then be opposition-panned for instant width and extra arrangement density. Tracks that offer this possibility will be those that include the same musical material performed more than once—not an uncommon occurrence in a lot of rhythm-section, riff, and hook parts. If you duplicate the track in question and use your audio editing tools to shuffle its repeated musical sections around, you can generate a version that combines with the original just like a separately overdubbed double-track. So let's say your song's chorus is eight bars long, but the chord pattern repeats twice within that. (If it doesn't, then it's your own silly fault for mixing prog rock.) It'll frequently be plain sailing to generate a double-track from a single strummed acoustic guitar part within that chorus, simply by copying the part and switching the order of the copy's two four-bar sections. When you mix the original and duplicate tracks together, they'll still sound like independent performances, because the original part will be playing bars 1 through 4 while the "double-track" plays bars 5 through 8, and vice versa.

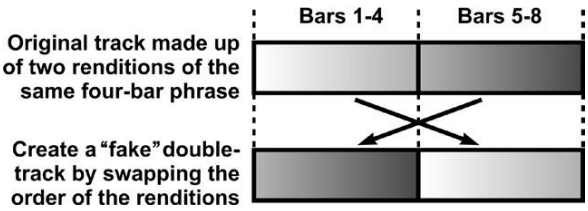

FIGURE 18.1
Generating a "fake" double-track with audio editing.

Where sections don't repeat quite as closely, though, you have to think a bit more laterally to do this trick. I regularly resort to pitch and time manipulation to square this particular

Mixing Secrets for the Small Studio.

circle, and we covered all the techniques you're likely to need for this back in Chapter 6. In other cases, I may build the double-track more as a kind of patchwork quilt by stealing individual chords and notes from all over the place to achieve the illusion of a fresh overdub, especially in those cases where a particular sound (overdriven electric guitar, say) throws its toys out of the pram if subjected to pitch shifting or time stretching. Some small-studio users try to create this effect by simply delaying a copy of a single track, but it's important to realize that this won't create nearly the same result and may cause serious phase-cancellation problems in mono. (There are some useful applications of that idea, though, which I'll get to in a moment.)

Another plug-in-free trick worth considering is adding some kind of wide-screen stereo background texture. A synth pad is a frequent choice here, but unpitched sounds can be equally handy. My stalwarts for this technique are things like tape hiss, vinyl noise, and the type of background room ambience used for film Foley work (often called "room tone" in the trade). The side effects of the pitched and unpitched additions will inevitably be a little different, though, so choose carefully. A pad will enrich pitched sustains, unpitched textures will to some extent improve perceived blend and obscure low-level details, and Foley ambiences may also clearly imply a sense of size and space. A particular advantage of using supplementary, but musically inessential, parts such as these to add stereo width is that mono compatibility rarely becomes an awkward issue. So what if your pad or room tone lose definition, or even disappear completely, in mono? Most mono listening environments are so lo-fi anyway that the typically subtle effects are unlikely to be missed.

18.2 ADJUSTING STEREO SPREAD

It's easy to adjust the spread of a collection of mono files in the stereo image using panning, but stereo recordings can be a bit trickier to deal with when it comes to adjusting stereo width. Granted, narrowing a stereo track isn't too tricky. All you do is adjust the individual panning of the left and right channels, as we saw in Section 8.2. Widening a stereo recording, on the other hand, is a bit like buying sensible undies: your best bet is M&S. Let me explain.

We're used to thinking of stereo recordings as being made up of one signal for the left-hand speaker and one signal for the right-hand speaker, but it turns out that you can also encode a stereo recording as two channels in a different way: one channel containing the mono elements at the center of the stereo image (the Middle signal) and one channel containing the elements that appear off-center (the Sides signal). As a sound moves off-center in a Middle and Sides (M&S) recording, its level reduces in the Middle signal and increases in the Sides signal, its polarity in the Sides signal indicating which side of the stereo field it happens to be on. You can represent an identical stereo signal in both left-right and M&S formats, and because each represents exactly the same information it's easy to convert between them.

FIGURE 18.2
A great little freeware M&S manipulation plug-in: Voxengo's MSED.

If M&S encoding doesn't make any difference to the sound of a stereo file, what's the point of it? Well, the reason it's so useful is that it allows us greater control over the mono and stereo aspects of a recording and in particular gives us the ability to increase the stereo width in a given recording by fading up the level of the Sides signal—push things far enough and you can even make the stereo soundstage appear to extend beyond the speakers in a stereo setup. However, it's crucial to remember that a stereo recording's Sides component isn't audible in mono, so if you start to heavily emphasize that component in the mix, the level drop when you move to mono may be quite dramatic—in extreme cases a sound could disappear entirely on some listening systems! So when increasing stereo width in this way, you should be sure to compare your mix balance in both stereo and mono, and it's also sensible to relegate such widening treatments to musically unimportant parts (such as rhythm-section fillers or synth pads) so that any level drops your widening does incur don't make a nonsense of the production.

The crucial thing to remember is that a stereo recording's Sides component isn't audible in mono, so if you heavily emphasize that in the mix, the level drop when you move to mono may be quite dramatic.

There's no shortage of stereo enhancer plug-ins that operate on M&S principles, although you can easily implement similar widening without any of them simply by duplicating your stereo track, panning the duplicate's left and right sides centrally (thereby summing it to mono), and then inverting the duplicate's polarity. As you add the duplicate into the mix, it will progressively phase-cancel more of the stereo recording's Middle component, thus increasing the relative level of the Sides component. Notwithstanding, I personally find it more convenient to use a separate M&S encoder/decoder, and my firm favorite for this is Voxengo's cross-platform MSED—not least because it's freeware, and I'm a cheapskate! This plug-in encodes a stereo left-right signal into M&S format within the plug-in, gives you separate level control over the Middle and Sides levels, and then reencodes the stereo signal to left-right format at its outputs. MSED also gives you the option of just encoding to or decoding from M&S, which offers additional possibilities if you wish to process the Middle or Sides signals of a stereo file independently using other plug-ins in your DAW—perhaps to EQ a central snare image in your stereo drum overhead recording without impacting as heavily on cymbal and room-sound elements, which appear closer to the edges of the stereo image.

Another application of M&S equalization is to provide frequency-selective width control, because any frequency you boost in the Sides signal will be widened, whereas any frequency you cut will be narrowed. However, this kind of frequency-selective operation is also often implemented by the equally viable approach of splitting the audible spectrum into bands and then applying separate M&S processing to each—exactly in the mold of a multiband dynamics processor. The reason frequency-selective stereo width processing is handy is

that it allows you to reduce mono compatibility problems with the minimum overall impact on the apparent stereo width. So, for example, if a phase misalignment between the low frequencies in the left and right channels of a stereo recording is causing a loss of bass when auditioning in mono, you could encode the stereo to M&S format and high-pass-filter the Sides signal to reduce the problem, without affecting the stereo impression delivered by the upper frequencies. In a similar vein, you might feel you can get away with boosting the very high frequencies of a stereo recording's Sides signal to give it an enhanced impression of width on high-quality stereo listening systems, knowing that the loss of this mono-incompatible high end will probably go unnoticed on lower-quality mono playback equipment.

18.3 STATIC ENHANCEMENTS

While M&S processing can do a lot to manipulate the stereo information that's already present in a signal, you can only really use it to enhance stereo width information that's already present. What if you want to add stereo width to mono tracks?

EQ-Based Widening

Simple EQ provides probably the most straightforward method. Here's the deal:

- Duplicate your mono track.
- Pan the original and duplicate tracks to opposite sides of the stereo image.
- Insert an equalizer into the original channel and make any old selection of cuts and boosts.
- Copy that equalizer setting to the duplicate channel, but with the gain settings of the equalizer inverted so that cuts become boosts, and boosts become cuts.

The effect of this setup will essentially be to pan different frequency regions of the sound to different positions in the stereo field, smearing the sound's spectrum across the panorama to produce a wider image. Although the advantage of this processing is that it's comparatively mono-friendly (the EQ changes pretty much just cancel each other out when the left and right channels are summed), it can produce rather an odd listening sensation in stereo, especially if you only use a few broad-band EQ changes. An EQ setting with lots of small peaks and troughs usually creates a more even feeling of frequency spread. This is why this technique is often associated with graphic EQ—if you set up an alternating pattern of boosts and cuts on the sliders for one channel, even a roadie should be able to invert that pattern by eye on the other channel. (As long as he understands long words like "alternating.") Nonetheless, I've never had particularly good results with basic EQ-based widening for important

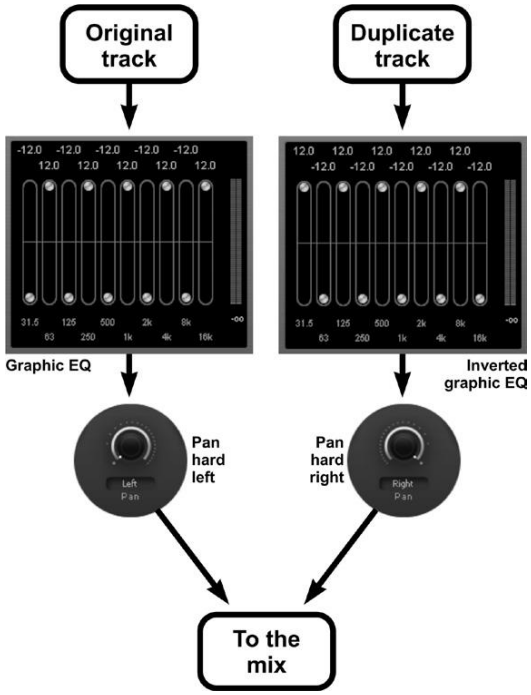

FIGURE 18.3
An EQ-based stereo widening effect.

music parts, because the way it destabilizes the phantom image seems to make instruments sound less solid in the mix—better to save it for incidental parts that suit a more diffuse and background role.

Pitch Shifting

Pitch shifting offers a different way to widen mono recordings. Follow the same track-duplication scheme as in the previous technique, but instead of inserting EQ to differentiate the channels, use small pitch shifts—try shifts of five cents (in other words five hundredths of a semitone) upward in the hard-left-panned channel and five cents downward in the hard-right-panned channel. Although such small pitch shifts aren't quite enough to make the track feel out of tune, they do fool the ear into thinking that the signals it's hearing from the two speakers are actually independent sound sources, so you get the impression of a wide-sounding stereo image.

One problem with this particular implementation, though, is that pitch shifters that sound any good are also quite CPU hungry, so you'll rapidly suffer a serious processor hit if you want to widen several tracks this way. For this reason, a much better real-world way of using this effect is to set up the pitch shifting as a send-return effect. That way you only need one pitch shifter for

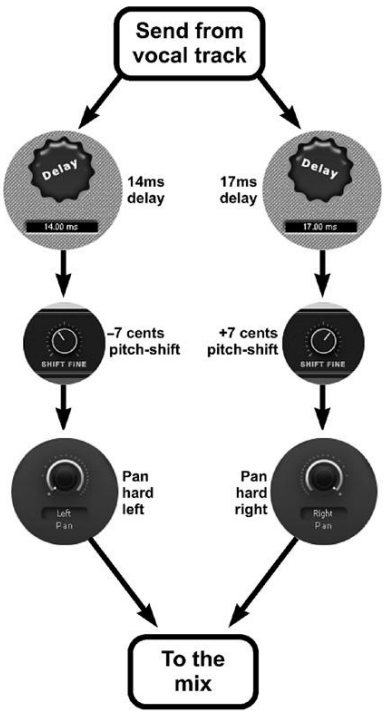

FIGURE 18.4
A classic stereo widening effect based around pitch-shifted delays, although every engineer has a favorite combination of pitch-shift amounts and delay times.

the whole project, and you can regulate how much it affects each instrument in the mix by adjusting individual track send levels. If your particular pitch shifter doesn't make a graceful transition into mono, you can also use M&S processing to reduce the level of Middle-signal component in the pitch shifter's return channel to reduce the nastiness for mono listeners, and you'll also usually want to add a short predelay (maybe 10ms or so) before the pitch shifter so that you avoid undesirable phase cancellation between wet and dry signals.

A Classic Pitch-Shifted Delay Patch

An extension of this pitch-shifting idea leads us to one of the oldest studio effects tricks in the book. Here's the way it works. If, instead of a single predelay before the pitch shifter, you actually use different delay values for each side of its stereo output (typically a few milliseconds less delay on the channel that is shifted downward), then the time-of-arrival difference between the two channels increases the perceived disparity between them and further increases the subjective wideness. All engineers seem to have their own preferences in terms of the exact pitch shift amounts and delay times, but so many people use this type of setup that it's got to be one of the all-time classic general-purpose mix treatments. Lead vocals are one of its prime applications, because not only does it widen them, making them appear larger in the stereo picture and decreasing the disparity between stereo and mono playback, but it also blends them ever so slightly into the track as a side effect, a little bit like the blend delay patch referred to in the previous chapter.

The biggest problem with this famous widener patch on lead vocals is that adding too much of it will make any band sound like Chicago, but even when you avoid that terrifying hazard, the pitch-shifted signal may still phase-cancel undesirably with the dry vocal sound. Although a little comb filtering may not matter on less important parts, on upfront lead vocals small changes will be more noticeable. One solution to this is to increase the predelay by another 20ms or so, but this means that you'll then begin to unmask the effect by letting it trail slightly behind the vocal. You get less comb filtering, sure, but you also can't get away with using as much widening before the pitch-shifted delay begins to sound obviously synthetic. Another tactic I personally prefer is to subtly finesse your delay times so that the comb filtering is rendered more benign—a millisecond either way can make a surprisingly big difference. I also frequently high-pass filter the effect return to some extent, which not only helps reduce the level of phase cancellation in the midrange (where it's usually at its most damaging in this instance) but also has the effect of supplying a little bit of high-frequency enhancement into the bargain.

Whichever way you set up this effect, though, you should make sure to check its mono-compatibility. One of the effect's strengths is that it's often fairly well-behaved in this respect, but if it isn't, then you can usually salvage the situation by massaging the pitch-shift or delay amounts while listening in mono.

Haas Delays

Another interesting effect for widening involves panning your dry track to one point in the stereo image and then panning a single echo to a different position, often on the other side of the field. The delay time here needs to be short enough that the echo can't be perceived separately from the dry signal and is therefore treated as if it were an acoustic reflection of the dry sound coming from another direction. In other words, it needs to be what is sometimes called a "Haas delay," after the German scientist who first explored this psychological effect experimentally. (Maybe it's just my schoolboy sense of humor, but I've always wished he'd been French.) If the delay signal can be perceived as a distinct echo in its own right, then you'll still get widening, but it'll sound much more noticeable as an artificial effect.

Typically a Haas delay will follow the dry sound within about 30ms, but this depends on the signal that's being delayed, so you always have to judge whether the delay is audible by ear—sharp transients can appear to flam at times well below 30ms, for example. You might be tempted simply to set the delay nice and short for safety's sake, but that brings problems of its own, because the further your delay time heads below 30ms, the more destructive the phase cancellation is likely to be between the wet and dry signals. You might not notice this if you've panned them hard left and right, but there'll be no escape in mono. The trick to getting best results here, then, is to try to set as high a delay time as possible without the echo breaking away perceptually from the dry signal. Some further processing of the delay return may help in this regard, along similar lines as when reducing the audibility of a blend reverb—high-frequency EQ cut and transient softening especially.

Typically a Haas delay will follow the dry sound within about 30ms, but this depends on the signal that's being delayed, so you always have to judge whether the delay is audible by ear—sharp transients can appear to flam at times well below 30ms.

For the widest stereo effect, you should pan the wet and dry sounds to opposite stereo extremes, but I find that most of the best applications for this Haas delay effect are where the wet and dry signals aren't pushed that far apart. Haas delays are great for gently expanding mono signals in sparser mixes where there's a lot of space to fill, for example. They're also useful where you want to avoid a big gap in the stereo field between hard-panned double-tracked parts—pan a Haas delay for each track somewhere closer to the center, or even overlap the delays slightly, to stretch

the instruments across as much space as you want to cover. When it comes to choosing a level for the Haas delay, it's tempting to use quite a lot, especially if you fall for the loudness bias by judging the impact of the effect just by switching the return on and off. However, Haas delays tend to rather homogenize instrument sounds they're applied to, smoothing over the internal details so that they sound a bit bland, and they also push sounds farther away from the listener, as you'd expect if you consider the effect's close relationship to blend delay/reverb. So tread a bit carefully here—listening in mono can really help clarify your decisions, because it takes the stereo widening out of the equation, thereby highlighting any potentially unwanted side effects.

18.4 MODULATED ENHANCEMENTS

So far we've looked only at static effects that can alter and enhance stereo width, but there are further options once you get into the realm of modulated treatments.

Auto-Panning and Rotary Speaker Emulation

Probably the simplest modulated stereo effect is auto-pan—an insert effect that sweeps a mono sound cyclically from one side of the stereo field to the other. Although this can be a nice effect on occasion, it's normally much too heavy handed for subtler stereo enhancements at mixdown and quickly becomes distracting for stereo listeners. However, there are still some useful widening patches to be had from an auto-panner if you keep the modulation depth low (so the panning isn't too wide), use a smooth modulation waveform such as sine or triangle (to avoid jerky stereo shifts), and push the modulation speed up into the 5 to 10Hz range.

FIGURE 18.5
Rotary-speaker emulations, such as the one within GSi's ORGANized trio plug-in, can work well as stereo wideners at mixdown.

Personally, I prefer to use auto-panning in a frequency-selective manner instead if possible, because it tends to feel less disconcerting if different frequency regions are panning in different ways—at least you don't get the whole sound charging rhythmically back and forth across the stereo field like some kind of audio tennis ball. Multiband auto-panners aren't exactly thick on the ground, though, so if you don't fancy building one from scratch, then an alternative is to find a Leslie-style rotary speaker emulation plug-in. At its simplest a plug-in like this will pan the low and high frequency ranges independently to simulate the speaker's two rotating horns. However, most models also modulate other aspects of the sound simultaneously, so you may wish to apply the simulation as a send effect to make the effects more subtle, perhaps with some predelay in the effect return to avoid comb-filtering problems. (My favorite freeware VST plug-in for rotary speaker emulation at mixdown is Mda's Leslie, mainly because I can set it up exclusively for frequency-selective auto-panning if I want, without any other processing side effects.)

Dynamic Tone and Pitch Changes

Although the static EQ widener I mentioned at the start of Section 18.3 has some applications, most of the time modulated tonal changes provide a more usable stereo enhancement. Chorusing, flanging, and phasing will all achieve this end when they're used as send effects, and while the exact mechanics of the way each works are slightly different (and not really essential to understand for mix purposes), their stereo widening tactics all essentially boil down to the same thing: modulating the tone of the left and right channels of the frequency response independently. Andy Wallace likes this kind of treatment on bass: "Not so much as 'Dig the flange on the bass!', because I kind of prefer that nobody even knows it's there…. I use that just to open things up a little bit so that everything is not kick, snare, bass right down the middle. Not that there's anything wrong with that, but sometimes it feels sonically a little more interesting to me."[1]

Only some implementations of these plug-ins will allow the individual stereo channels to be processed separately like this, and others have this facility switchable, so your first setup step will be to choose an appropriate plug-in and preset. Beyond that I normally feel that the most successful widening effects for mixdown arise when the effect's Feedback and Depth controls are kept low, although the strength of the widening will, of course, depend on the send level to the effect from the dry track.

Modulated pitch shifting, better known as vibrato, can yield widening effects if you can actually find a plug-in that offers it (audio vibrato processors aren't that common) and then persuade it to treat the stereo channels independently. If you have this option available, keep the vibrato's Depth control low again to avoid tuning vagueness creeping in, and consider introducing some predelay into the effect return to reduce unwanted tonal colorations.

Whatever modulation effect you use, it is still vital to check how your processing sounds in mono. One of the good things about dynamic widening effects

FIGURE 18.6
If you can find a modulation plug-in that allows its left and right channels to operate out of phase, then that may work well as a stereo widener. Here you can see a couple of freeware examples. GVST's GChorus provides variable phase for its right channel, whereas Kjaerhus Audio's Classic Chorus offers a simple Spread button.

is that the damage caused by undesirable phase cancellations will tend to be softened by virtue of the fact that the modulation will be varying them the whole time. Still, you can't know how your mono-compatibility stacks up without using your ears.

18.5 CHOICES, CHOICES...

I've covered a lot of options for stereo widening here, and you might be forgiven for wondering why you actually need more than one. The reason there are so many approaches is that each one creates a different subjective sound, exhibits different side effects, and affects mono-compatibility in different ways. It's only by trying them all out and developing a memory of what they sound like that you can choose the best one for any given widening task.

CUT TO THE CHASE

- Stereo width is at least as important to modern record production as exact stereo positioning, so it's important to have a selection of stereo width-manipulation techniques at your disposal. All of them will change the way stereo and mono listening affects your mix balance, so you should frequently check mono-compatibility as you're working with them.
- Generating "fake" double-tracks via audio editing is a very effective strategy for filling out many small-studio productions, both in terms of texture and stereo width. Mixing in a delayed copy of a track doesn't produce the same effect. Pads, tape hiss, vinyl noise, and Foley room tones can all add stereo width too, alongside note-sustain, blend, and size side effects.
- M&S stereo encoding/decoding gives you a lot of control over the image width of stereo signals, particularly when applied in a frequency-selective fashion, but it can also lead to serious mono-compatibility problems if pushed too far.

- Static EQ, pitch shift, and delay effects can all deliver stereo widening artifacts, as can modulation processes such as auto-panning, rotary speaker emulation, chorusing, flanging, phasing, and vibrato—assuming that your particular DAW and plug-ins allow you to process the two sides of a stereo signal independently.

Assignment

- Do a survey of the plug-ins on your DAW system, and try to find at least one of each of the following: auto-panner, rotary speaker simulator, stereo chorus, stereo flanger, stereo phaser, and stereo pitch-shifter.
- If you can see any fruitful avenues for editing together "fake" double-tracks in your own example mix, then try some out to see whether they might fill out the arrangement texture and stereo field in a beneficial way.
- Return to your mix and experiment with a good selection of the different stereo adjustment effects covered in this chapter. Try at least M&S-format rebalancing, static EQ-based widening, the classic pitch-shift delay patch, Haas delays, and a couple of different modulated stereo treatments. You shouldn't need to use all of those on every mix, but having a feel for the sonic characteristics and side effects of them all will better prepare you to choose the most suitable option for any given task.

http://www.cambridge-mt.com/ms-ch18.htm

This page intentionally left blank

CHAPTER 19

Buss Compression, Automation, and Endgame

You should now have the most important section of your mix balanced and sweetened to taste, which brings you to within striking distance of a complete first draft. This chapter explains the steps you need to follow to go the remaining distance.

19.1 BUSS COMPRESSION

At this point in the mix, it makes sense to start considering any buss compression—in other words, compression of the whole mix via a plug-in on the main mix buss. This is considered desirable within many music styles, for several potential reasons:

- It introduces some level interaction between different parts in your mix by ducking quiet signals slightly in response to louder ones, and in doing so it gives a feeling that the mix coheres better. It's common to hear engineers talk of buss compression "gluing the mix together" on account of this.
- It can create pumping effects that add a subjective sense of loudness and aggression, which suits certain styles.
- If a characterful compressor is used, then it may subjectively enhance some aspects of the mix—for example, by adding subtle distortion artifacts at certain frequencies.
- Some engineers feel that it reduces the negative side effects of down-the-line transmission compression. Here's Andy Wallace, for example: "A long time ago, I learned that [radio compression] was seriously changing how the low end sounded and the balance…. That's when I really started experimenting with a substantial amount of stereo compression. And I found that if I had something compressed ahead of time and was happy with the sound of it, the additional compression from the radio station had less effect."[1]
- It evens out the dynamics of the entire mix signal, bringing out internal mix details and increasing the production's average levels so that it appears louder at a given metered peak level.

Mixing Secrets for the Small Studio.

Of these, the last isn't actually a very good reason to apply buss compression during mixdown, because loudness processing is much better carried out at the mastering stage when the mix is complete. It's also important to remove loudness bias from your decision-making process when applying buss compression so that you're not misled into overprocessing. In other words, adjust your compressor's Makeup Gain control to compensate as best you can for any loudness increase, so that bypassing the compressor doesn't boost the subjective volume.

As far as which model and setting of compressor to use, my own research suggests that opinions in the professional community are many and varied, so buss compression isn't just some kind of preset you can slap over your mix without thinking. When a compressor is acting on a complex full-range mix signal, the sonic differences between compressor models and the effects of small adjustments to fine parameters begin to make much more profound sonic differences, and subjective judgments are often as much about stylistic conventions and the engineer's personal preferences as anything else. However, that insight isn't going to help you much in finding a setting that suits your particular mix, so let me at least suggest some of the practical factors you may need to consider.

Buss compression simply for the purposes of "mix glue" rarely uses more than about 2 to 3dB of gain reduction at most, and it often relies on a slow attack time and automatic release time to keep transient smoothing and pumping artifacts to a minimum. Ratios over 2:1 are unlikely to be useful in these cases, and you may discover that ratios as low as 1.1:1 may be more appropriate, depending on the musical style. If you're looking for more aggressive pumping compression, however, then you might even see 8dB of gain reduction on peaks, with faster attack and release times and possibly also higher ratios. Lee DeCarlo also recommends one important refinement: "I'll play with the release and the attack times until I can actually make that limiter pump in time with the music."[2]

But no matter how much your band look like extras from *The Lord of the Rings*, the common thread with any type of buss compression is that you shouldn't compress the mix buss any further than the point where unwanted processing side effects become overbearing. Much of the art of successful buss processing therefore lies in working out how to rein in unwanted side effects as far as possible, so that you can maximize the desirable aspects of the sound. We looked at these side effects in Chapters 9 and 10, but here's a quick refresher of the main offenders, with suggestions of how you might deal with them:

- *Loss of attack on prominent transients, such as kick and snare drums.* Suggested remedies: reduce gain reduction; increase attack time; use parallel compression (perhaps via a compressor with a wet/dry mix control); reduce ratio; adjust knee softness; switch level-detection mode from peak to average (RMS); select a different compressor design.
- *Excessive gain pumping.* Suggested remedies: reduce gain reduction; adjust release time or switch to automatic mode; apply low-frequency EQ cut in

the level-detection side-chain; use parallel compression; select a different compressor design.

- *Unwanted distortion.* Suggested remedies: lengthen attack and release times; adjust knee softness; select a different compressor design.
- *Loss of weight on instruments with low-end transients, such as kick drums.* Suggested remedies: reduce gain reduction; lengthen attack time; apply low-frequency EQ cut in the level-detection side-chain; adjust knee softness; select a different compressor design.
- *Undesirable alterations in the subjective mix balance.* Suggested remedies: reduce gain reduction; lengthen attack and/or release times; adjust knee softness; select a different compressor design.
- *Unappealing tonal changes to the mix.* Suggested remedies: select a different compressor design.

If you're having trouble hearing the side effects of your buss compression, the best thing is to drive the compressor a bit too hard to start with. Once you've got a handle on what you're listening for, you can reduce the amount of compression to more sensible levels. Sometimes a few of these side effects may be impossible to dispel entirely with adjustments to the compression alone, without losing the desired degree of mix glue, pumping, or compression character. In this event, an additional remedy can be to adjust the mix itself to compensate for the compressor's less delectable side effects. You might enhance the kick and snare transients with additional processing to keep them crisp despite the attack-dulling effects of heavy pumping compression on the mix buss. Or you might tweak a few channel faders if some instruments (or particularly effects returns) have been knocked out of balance by the buss processing. Both of these examples are common practice, and it can be next to impossible to match the sounds of competing commercial productions in some styles unless you make use of these kinds of tactics yourself. By the same token, though, you will need to keep readjusting your buss compression settings as the full mix continues to develop, in order to keep the plug-in operating in its most magical-sounding zone. It's because of the interactions between individual channel processing and buss compression that it's essential to get any mix-buss compressor involved during the mixdown process itself, rather than leaving it until the mastering stage as you would most other full-mix processing.

No matter how much your band look like extras from *The Lord of the Rings*, you shouldn't compress the mix buss any further than the point where unwanted processing side effects become overbearing.

If you're having problems with gain pumping, why not use multiband compression instead? Although in theory there shouldn't be a problem with this technique, because the same interaction principles apply, the reality is that multiband buss compression is almost always best left to the mastering engineer—for practical reasons as much as anything. Multiband processing is

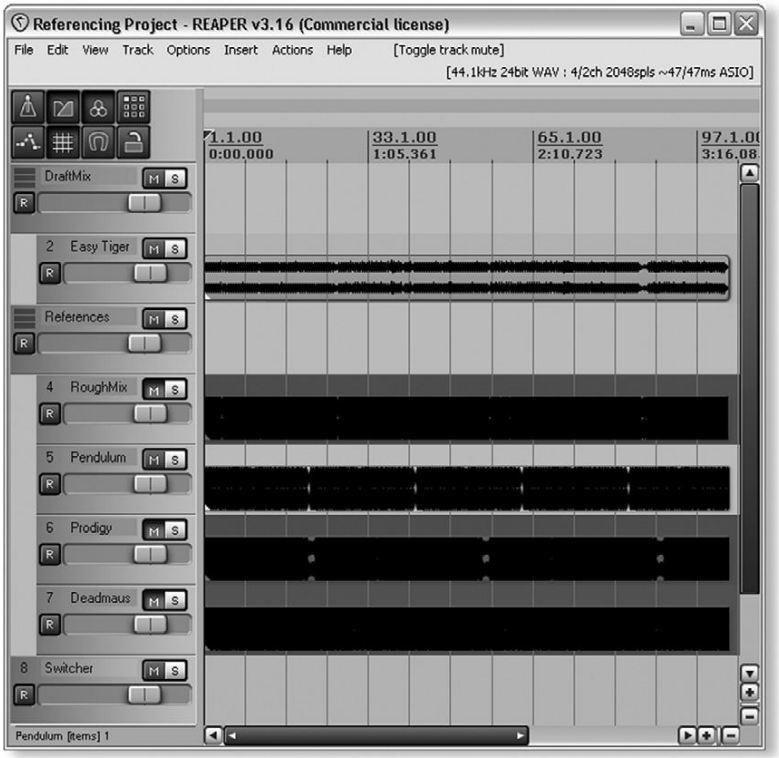

FIGURE 19.1
For referencing purposes, it's good practice to export your mix into a separate DAW session so that you can easily compare it to a selection of commercial productions.

just too complicated and finely tuned to manage alongside all the other aspects of the mix, and everyone I've ever seen using it on their mix buss has fallen into at least one of the following traps: they put insufficient effort into getting the balance right, because the multiband compressor's automatic frequency adjustments make spectral balance problems in the mix more difficult to spot and remedy, or they undercook their blend effects and then overcompress the buss in an attempt to compensate with additional "mix glue," or they overprocess individual tracks to compensate for dynamic frequency-response side effects of the buss treatment. In short, full-band compression is more than enough to worry about while mixing, so save yourself some grief and leave it at that.

19.2 LOUDNESS MATCHING FOR REFERENCING

In an ideal world, once you've got your buss compression up and running and you've rebalanced individual tracks as necessary in response to it, you should have arrived at the final sound for your most important mix section. But, sadly, the world isn't ideal! The cold, hard, truth of the matter is that at least a handful

of your mix decisions are probably less well-judged than the contents of Noddy Holder's tie rack. There's no shame in that, though, because the only reliable way to iron out the remaining wrinkles is to line up your mix against the commercial competition, and it's only now you've got a fully mixed song section that you're properly placed to engage in this comparison process. So export the section of the production you've mixed so far as a stereo audio file, and import it into a fresh DAW session along with some appropriate reference material, perhaps made up of specific target productions suggested by the client, as well as some tracks selected from your personal reference library (as discussed in Chapter 4).

As I mentioned in Section 4.2, the most important things to do when making these comparisons are to switch between the mixes instantaneously, so that your hearing system has no time to adapt itself to differences in balance and tone, and to adjust the subjective levels (not the metered levels!), so that loudness bias doesn't interfere with your quality decisions. However, in styles of music that use loudness processing at the mastering stage (which is pretty much everything these days), you might feel you want to apply similar processing on your own unmastered mix section in order to make truly meaningful comparisons. Loudness maximization, like all other processing, has side effects, so you need to be aware if these might require compensatory tweaks to aspects of the mix itself.

Basic Concepts of Loudness Maximization

The basic purpose of loudness maximization is to cram the maximum subjective level onto fixed-headroom digital delivery media such as CDs and MP3 downloads, but without unacceptably compromising other qualities of the production. There are lots of ways to achieve extra apparent loudness without increasing peak signal levels, each of which has its own cost in terms of potentially unwanted processing artifacts, so choosing between them (or indeed combining them) depends on how much loudness you want and what kinds of side effects you're willing to tolerate. These decisions must reflect the wishes of the client, the expectations of the target market, and the way the mix in question responds to the different processing options. Demacio Castellon suggests, "If you're making a jazz record you're not going to make it as loud as a pop record. You're not going to make a blues record louder than a heavy metal record."[3] While I can't make the choices for you (your client, target genre, and mix will be unique, after all), what I will do nonetheless is describe some common loudness processing options so that you're aware of the main side effects in each case. That way you should be able to come to reasonably well-informed conclusions when working with loudness maximization processing for reference purposes.

But before I get into specifics, I want to stress a crucial general principle that should be at the front of your mind when flirting with loudness processing: it's the side effects you should be concentrating on, not the loudness hike. This has two main ramifications, the first of which is that you should never, under

pain of a visit from the Audio Engineering Society's ninja death squad, evaluate the success of a loudness processor by comparing it to a subjectively quieter unprocessed signal. If you don't match the subjective loudness of the processed and unprocessed sounds for comparison purposes, then the loudness bias will render any subjective quality evaluations worthless. Most loudness processors make it ridiculously easy for you to mug yourself in this way, so you have to be super vigilant. My preferred working method is to duplicate my mix to two separate stereo mixer channels, but to process only one of them for loudness purposes. I can then switch between the two for comparison and adjust the faders for the best subjective loudness match.

A crucial general principle that should be at the front of your mind when flirting with loudness processing is this: it's the side effects you should be concentrating on, not the loudness hike.

The second ramification of focusing on subtle differences in processing side effects is that it demands that you push your monitoring technique to the limits. Some side effects aren't necessarily very audible on one type of listening system, for example, so you can only make well-rounded judgments about the relative merits of different loudness enhancements based on information from all quarters. Other side effects may only be noticeable at certain listening volumes, in certain places in your control room, or on a spectrum analyzer. There are no shortcuts to achieving sensible loudness processing in this regard, so consider yourself warned.

Suggested Processing Strategies

With those stern provisos out of the way, let's have a look at some actual processing strategies and their side effects:

- *Full-band "top-down" squeeze.* This is where you attempt to gently squeeze a large section of a signal's dynamic range, using very low-ratio compression operating above a low threshold value. As long as gain reduction is kept within 3dB or so, it's usually possible to keep pumping artifacts relatively benign by adjusting attack/release times by ear. Potential side effects include undue emphasis of low-level details such as background sounds, reverb/delay effects, and incidental noises; unwanted overall level increases during sections with sparser arrangement; and reduction in transient definition.
- *Full-band "bottom-up" squeeze.* A similar approach as in the top-down squeeze, except that the dynamic range below the compressor's threshold is targeted. There are specialist processors for this purpose (called upward compressors or de-expanders), but you can also achieve a similar effect using subtle full-band parallel compression. Potential side effects are similar to those of top-down squeeze—transient definition tends to suffer less, but you get less peak control and an increased likelihood of unwanted overall level increases during sparser arrangement sections.

FIGURE 19.2
This screenshot of Flux's Pure Compressor II plug-in shows a typical gentle setting for full-band top-down squeeze.

- *Full-band limiting.* This is usually a much faster-acting gain reduction designed to stop signal peaks dead, while leaving the remainder of the dynamic range comparatively unscathed. In some types of music, an intelligently designed peak limiter can achieve a peak reduction of several decibels before audible problems start arising. Potential side effects include pumping, bass distortion, softening of transients, and reduction of apparent drum levels.
- *Multiband compression/limiting.* Using a multiband configuration for any of the three dynamics processes I've just described gives you additional scope for dynamic-range reduction before pumping artifacts impinge on your enjoyment of the music. However, pumping is only one of the potential side effects of compression-based loudness enhancement, and if you use multiband processing to increase the amount of gain reduction, then you can easily end up with greater side effects in terms of mix-balance alteration and transient softening. Any multiband approach also adds further problems of its own: because the amount of gain-reduction in each frequency band will depend on the overall level of frequencies within it, changes in the spectral content of the mix (perhaps as a result of alterations in the instrumentation) can trigger unwanted changes in the overall mix tonality as the compression/limiting adjusts its gain-reduction—it can be almost as if some maniac were randomly tweaking an EQ over your whole mix in real time. There are also engineers who feel that multiband gain reduction

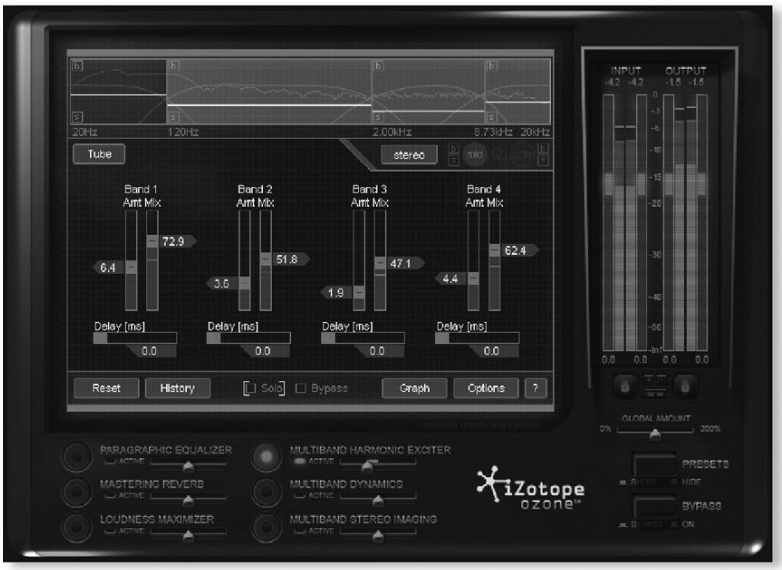

FIGURE 19.3
Some mastering-style processors (such as iZotope Ozone, pictured here) offer subtle multiband distortion treatments that can boost the subjective loudness of a mix.

drains the life out of a production by ironing out the tonal contrasts between its musical sections.

- *Subtle distortion.* By adding distortion harmonics to a mixed signal, you can increase its harmonic density and apparent loudness with very little increase in its peak signal levels. Subtle valve, tape, and transformer distortions are all options here, and if you need greater control over the exact nature of the harmonics additions, then parallel or frequency-selective configurations may be appropriate. Potential side effects include fatiguing tonal harshness, emphasized vocal sibilance, increase in the apparent level of treble percussion instruments in the balance, veiling of midrange details, and unwanted overall changes to the mix tonality.
- *Clipping.* Yes, clipping—whether it's straight digital-style flat topping of the waveform peaks or some kind of modeled analog saturation that rounds them off more smoothly. You can find any number of textbooks that threaten hellfire and brimstone should you dare abuse your full mix in this way, but such admonitions are at odds with widespread commercial practice—examine the waveforms of the top 40 singles any week of the year and you'll see clipping in abundance. The advantage of clipping as I see it is that it doesn't seem to affect the subjective attack or balance of prominent drum parts as much as peak limiting, so it tends to suit styles with hard-hitting rhythm parts. Potential side effects include subjective tonal change of clipped peaks and unwanted distortion on more steady-state signal waveforms (although this may be disguised to some extent by distorted instruments, such as electric guitars, within the music itself).

Of course, each of these strategies will suit different productions, and each will only give you a certain range of loudness enhancement before it begins to reveal its weaknesses, so it's not unusual to find that several loudness processes in combination are required to push your mix as loud as your references. You may also quickly discover that it's still impossible to push your mix to the target loudness without unconscionably heinous side effects, and this can be a clue that some aspect of your mix itself needs to be tweaked to allow the loudness enhancement to work more effectively. For example, drum peaks commonly find themselves bludgeoned down into the mix balance by extreme loudness processing, so you may need to boost the apparent power of these instruments by increasing their sustain instead.

19.3 REFERENCING CHECKLIST

Once you've managed to match the loudness of your draft mix section fairly well against your reference material, you need to take a deep breath, leave your ego at the door, and submit every aspect of your mix to merciless side-by-side comparison. Give all your monitoring systems a good workout, check at a range of different listening volumes, and don't stop until you've winkled out every last insight. Remember that this mix section should be the high point of the production, so it should shake the listener's socks off one way or another. Here's a quick rundown of issues I usually make a conscious point of assessing:

- *How does the overall mix tonality compare?* This is something that can easily be out of line, because it's difficult to stop your ears from adapting to a skewed tonality while you're stuck into detailed mix activities. If the tone feels wrong, then insert an equalizer over your mix to see if you can correct the frequency imbalance. You can subsequently apply this same EQ setting on your mix project's master buss or rebalance the frequency response of individual instruments with reference to it.
- *How does the balance compare?* Concentrate primarily on the most important instruments here, because they usually offer the least leeway within any given musical genre. I usually make a point of checking at least kick, snare, bass, and lead vocals, but the more parts of the mix you direct your attention toward specifically, the more representative you'll be able to make your own balance. Don't just listen for overall levels, but also consider how much dynamic range each instrument has.
- *How does each instrument's tone compare?* Again, concentrate primarily on the most important instruments. Also, if the tone of any instrument in your mix appears to be less appealing, then ask yourself whether there's actually anything you can do about that without upsetting the balance.
- *How does the use of reverb and delay effects compare?* Listen for how well the different instruments blend with each other and how any different size reverbs have been used. Give a thought to overall reverb/delay lengths and levels as well.

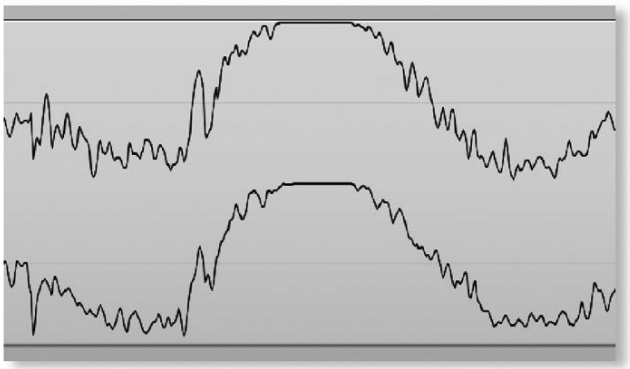

FIGURE 19.4
A lot of high-profile albums use clipping as a means to increase subjective loudness—the burst of more than 100 consecutive clipped samples shown here (from the Kaiser Chiefs' hit "Never Miss a Beat") is by no means unusual these days.

- *How does the stereo image compare?* This is partly a question of the apparent width and placement of each individual instrument, but you should also consider how wide the stereo image is overall in different regions of the frequency spectrum.

If you've selected good reference material, you're using all your monitoring systems effectively, and you're being brutally honest with yourself, then you should end up with a mix adjustment "to do" list as long as your arm. You may even feel like you're back at square one, but try not to be daunted. This is the make-or-break point of any mix as far as improving your skills is concerned, so if you can keep your nerve you'll be putting yourself on the fastest track to commercial results. "You're going to make mistakes," reassures Humberto Gatica. "The important thing is to learn from them, to tell yourself 'Oh my God, that sounds terrible, I will never do that again.' … You correct yourself and then you balance it so the good overrules the bad. And then you go on from there."[4] I can't stress enough that the difference between small-studio engineers who can mix and those who can't is that the latter fight shy of proper referencing because they can't handle a bit of disappointment in the short term. What are you: man or mouse? You can read all you like about processing techniques, but it's the referencing process that actually teaches you how to mix.

So have a sushi break, catch up with the latest Simpsons, and then steel yourself to head back to your mix project and deal with those revisions. When you're done, export that section of mixdown and run through the whole referencing process again, keeping all your different monitors in constant rotation. And keep repeating that process as long as you can stand it. "The idea is to make it work on all systems," says Andy Johns. "This is a lot of work, but it's

the only way to go."[5] Worked your way through a dozen different versions? Hah! You're clearly not taking things seriously enough.

19.4 AUTOMATION FOR LONG-TERM MIX DYNAMICS

The goal of the recurring nightmare that is the referencing process is to arrive at a version of your climactic mix section that is not just earth-shatteringly magnificent on its own terms but also stands shoulder to shoulder with the market leaders. If your references are any good, you'll probably find that you never quite reach this goal, but that's fine—you're trying to clear an incredibly high bar, so even if you don't quite make it, you should at least end up with something that is solidly of a commercial standard. Your job now is to fill in the remainder of the mix sections. I described the general principle behind this step in Section 8.1, but there are a few additional points to raise based on what we've covered since then.

The way I usually set about building up the rest of the musical structure is by muting all the currently active tracks, looping the new musical section, and then reintroducing the tracks in the appropriate rank order. Whenever a track needs adjustment to fit its new context, I do one of two things: either I mult that section to a different track, or I engage my DAW's mixer automation facilities to handle the change. Bob Clearmountain uses a similar section-based approach: "In the '70s, when we were recording on analog with no automation… we would mix sections and then edit them all together, and I tend to work pretty much the same way on the computer…. I'll try to get each section sounding right until I have a complete pass that is a pretty good basic mix."[6] Mults are best for situations where a track's processing and effects need to change in a variety of different ways simultaneously, because it's easier to deal with lots of changes from the mixer controls than by tweaking dozens of streams of automation data. Where it's just a couple of parameters, though, automation is usually a more efficient solution and will save the CPU overhead of duplicating plug-ins.

The goal of the recurring nightmare that is the referencing process is to arrive at a version of your climactic mix section that is not just earth-shatteringly magnificent on its own terms but also stands shoulder to shoulder with the market leaders.

Automation of practically any mixer or effect parameter is now a standard feature of almost all DAW systems, and if it's not then sling that software into the trash and move to a different platform! Seriously, it's almost impossible to compete with commercial releases in any musical genre without the use of mix automation, so it should be considered a prerequisite. "Fader moves are even more important than EQ," says Chris Lord-Alge. "No matter what song you have, you need to help it build."[7]

BUSS EQ

Because our ears are so good at adjusting to differences in tonality, it's common to discover while referencing that the overall mix tone needs tweaking. Although you could deal with this on a track-by-track basis, it's more straightforward to just add an extra EQ plug-in over the main mix channel—put it after any buss compression, though, so that it doesn't interfere with your finely honed settings there. The big advantage is that you can use good-quality EQ for the purpose, and this is particularly important in the common situation where you want to increase brightness. Spike Stent explains the rationale here: "Basically what I do with every mix is put a GML EQ across the stereo buss of the SSL, and just lift the top end, above 10K, so I won't have to use as much high end from the SSL. The mid and high end of the SSL is a little harsh and cold. So I turn that down a bit, and then lift the high end with the Massenburg to make the mix sound warmer and add sheen."[8] Although your choice probably won't be between an SSL console and a GML EQ, the same reasoning applies when working on DAW systems, because adding high end in the digital domain with lots of low-CPU channel plug-ins tends to introduce more harshness than when you use a single high-CPU plug-in over the master outputs. "Digital EQ doesn't handle high frequencies very well," says Jacquire King.[9]

By using mults and automation, you can slowly assemble the overall production piece by piece, and as long as you pay enough attention to the order in which you tackle the sections, your long-term mix dynamics should make fairly good sense once all the gaps are filled. If the arrangement isn't particularly inspired, however, then you may need to reinforce the long-term dynamics, in which case here are a few tips to get you started.

One of the most hackneyed techniques (but no less useful) is to contrast the levels of reverb/delay between the verses and the choruses. Mark Endert says, "You can get a lot of contour out of a track that lacks dynamics just by changing the acoustic space around it."[10] The most common configuration is for the verses to be drier and closer, where the choruses are bigger and more live sounding. "The reverb that would sound plenty wet in a sparse part of the track will be too dry as the track builds dynamically," says Mick Guzauski. "So, often, I'll ride the reverb and delay as the track builds."[11] So when that first chorus hits, push up the room mics, or fade up the return from the parallel compressor on the overheads, or increase some of the sends to the size reverb, or switch on an epic tempo delay. (When you reach the final chorus, you might go the whole hog and do all of those things at once!) Once you've got chorus 1 really opening out, you then have the option to catch the listener's attention at the start of verse 2 by abruptly cutting the reverb and delay tails dead just before the lead singer starts their first line. Bear in mind, though, that it often makes sense for verse 2 to be slightly less dry than verse 1, given that chorus 1 has by that point introduced the idea of a more reverberant sound.

Also firmly in the "cliché for a reason" category is a similar approach toward stereo width. Widening out your choruses, or perhaps just packing their stereo

picture more densely, is a time-honored crowd-pleaser. This is where having stereo enhancement send effects can really pay off, because you can easily ride their levels as in the previous reverb situation. Introducing opposition-panned "fake" double-tracks, Haas delays, or pads only for the choruses would be alternative options to weigh up.

If one of the instruments in a packed-out chorus arrangement also appears in less crowded conditions elsewhere, don't forget to reassess its high-pass filter frequency within the new setting, particularly when the instrument in question is something acoustic. Piano parts are a regular case in point here. They often require fairly severe low-frequency cuts in the context of a busy mix, but that same EQ will make the instrument sound like a tap-dancing vole if it ever moves further into the limelight.

Pads regularly require automation adjustments if you want them to consistently enhance the mix without anyone noticing. A pad that plays across several different musical sections will almost certainly need at least level automation to retain the perfect enhancement/inaudibility equilibrium, although you may also find that the EQ has to change too if the synth is masked much more heavily in some arrangement textures than in others.

Automation can be a great boon when you've got to create any kind of buildup, whether that's a short ramp-up into a chorus or the slow development of an extended outro workout. There are two important points to realize when mixing buildups. First, the peak of any crescendo must inevitably be fixed to fit in with the production's long-term dynamics, so the only way to

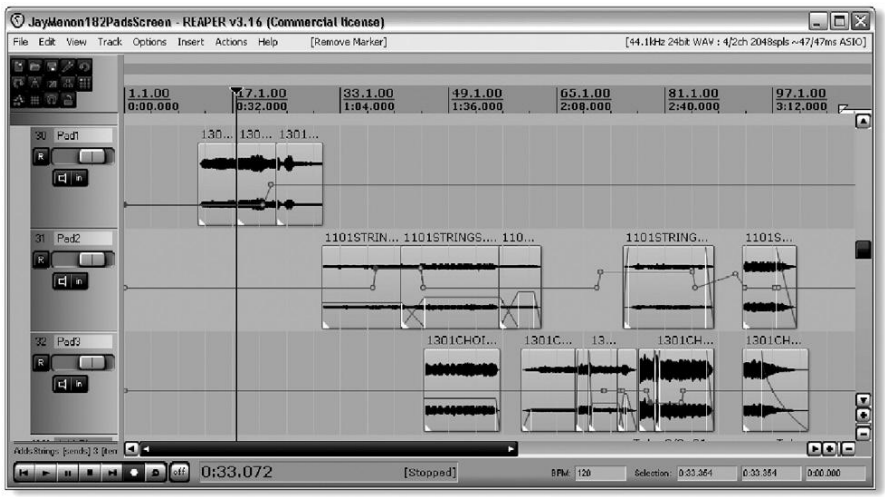

FIGURE 19.5
If you're using synth pads of any kind in your mix, you'll probably need to automate their levels so that they adequately fill out the lush mix moments, but don't also become too audible in sparser sections of the arrangement.

create an impressive buildup is to drop the overall excitement level as far down as you can in anticipation—the less you drop the energy levels, the less room you leave yourself for maneuver. The second crucial point is that you should try to rely as little as possible on sheer level increases to implement any buildup, because these won't survive downstream loudness processing particularly well. It's much better to intensify factors like arrangement complexity, instrument timbre, effects levels, and stereo width to create the majority of the impression, as these illusions will tend to survive better than level increases in real-world listening situations.

19.5 DETAILED RIDES

If you mix each section of your production according to the principles I've discussed so far, you should eventually end up with a complete mix that balances sensibly throughout, while at the same time appearing to progress in a logical and musical way from start to finish. However, it will still need one final ingredient to bring it up to a professional level: detailed rides. These short-term automation-data wiggles are what take a mix beyond what plug-ins can achieve, adding in the essential human element of intelligent, dynamic musicality.

Although I briefly mentioned the idea of detailed rides when discussing lead vocals back in Chapter 9, I've otherwise deliberately glossed over their role until now because they're something of a two-edged sword. On the one hand, automation is the most powerful and musical mix process you can possibly apply, and it is capable of matching and exceeding the capabilities of even the most advanced dynamics processors. An automated fader will produce more natural results than any compressor, for example, because it's got a human being as its level-detection side chain and gain-reduction circuit. An automated EQ will remove sibilance more sensitively than any automatic de-esser, because your ears will be much better than any DSP code at judging the exact nature of the frequency-selective gain reduction required in each instance. An automated reverb send can respond to arrangement changes in a way that no plug-in can manage.

But on the other hand, detailed rides are labor intensive too, and it can be difficult to implement broad-brush changes to automation data after it's been created. It might take you a couple of hours to compress a single bass line using automation, whereas compression will do most of the job for you in minutes and will let you readjust the attack time easily at a later date. There's no sense in de-essing everything manually if an automated de-esser will do the bulk of the work for you in a fraction of the time. You can achieve useful improvements to your reverb levels by using a keyed ducker in the reverb return, it'll save you ages automating a bunch of sends. Although there's nothing stopping you using automation to construct a perfectly good mix without the help of a single dynamics plug-in, you'd have be an utter fruitcake to work that way voluntarily.

The upshot is that it's not usually a great idea to get rides involved too early in the mixing process if you can avoid it. The closer you can get to a great mix without them, the less time you'll have to spend fiddling with little graphical curves in your DAW system and the more productive your limited mixing time is likely to be. Reserve rides only for those situations where multing and dynamics processing fall short, then you shouldn't waste too much time. As such, rides usually serve one of three main functions: they troubleshoot mix problems that mix processors are unable to isolate effectively, they improve the subjective mix balance beyond what processing can achieve, and they direct the listener's attention toward interesting arrangement features via momentary balance changes. Let's look at each of these roles in turn.

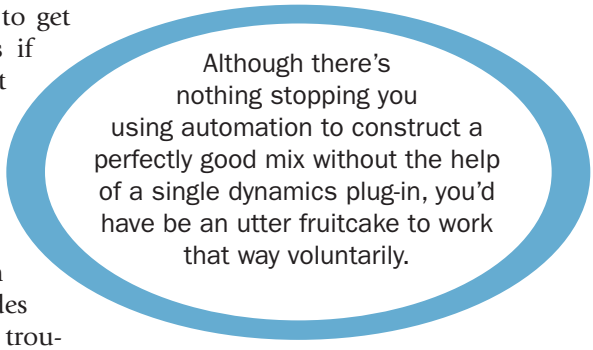

Intelligent Troubleshooting

Time-domain processing is all about prediction—your job is to listen to the moment-by-moment variations in a track, build a kind of generalized mental picture of what's going on, and then forecast the kind of gain-change processing that might be able to fit them all into the mix. Unfortunately, though, music recordings aren't always predictable. For example, your acoustic guitarist might normally be very good at keeping any fret noise under control, but just before the second chorus he accidentally plants a real humdinger. Dynamics processing won't supply the answer here, because the kind of drastic processing you'd need to target that lone fret squeak would also mash up the rest of the performance. Assuming that there's no option to patch over that transition with audio copied from elsewhere in the song, you could mult this squeak to a separate track to sort out the problem, but it's usually more elegant to deal with this type of isolated blemish by using automation instead.

There are lots of other common applications for these kinds of rides:

- Dipping vocal breath noises.
- Adding attack to a fluffed guitar note by fading up the initial transient.
- Rebalancing individual bass notes. Comments Mike Shipley, "I find that if the player's pickup is substandard, one note will 'out-boom' another, and though I'll first try to fix it with a dynamic EQ, it rarely works. So I'll just automate the track, dig in and find the frequency, and smooth the bass out."[12]
- Zapping momentary feedback howlarounds on live recordings.
- Catching overprominent "s" or "f" consonants that manage to bypass your de-esser. Steve Hodge states, "I routinely de-ess with automation now by [ducking] the level of the esses down without affecting the rest of the word. It can be that precise."[13]

- Reducing the low-frequency thump when a mic stand has been knocked during recording.
- Compensating for changes to recording settings accidentally made during a take: a performer might have moved position in the height of artistic inspiration, or the engineer might just have nodded off onto some faders!

FIXING BAD COMPRESSION WITH AUTOMATION

Detailed automation provides a solution to one of the trickiest mix-salvage jobs. Tom Lord-Alge explains: "Often I want to uncompress things that are on the tape, when they are overcompressed or compressed in a nonmusical way. But I fix that by compressing even more and then creating new dynamics by using the faders. Automation comes in handy for that."[14]

Dave Pensado has his own slant on this tactic for vocals:
"I go into Pro Tools and I type in the level on every syllable… and I automate the EQ on every syllable… On a four-minute song, to fix a lead vocal with just the most horrible compression takes me about two hours."[15]

Perfecting the Mix Balance

The problem with most dynamics processing for balance purposes is that it's only concerned with the track it's processing, so it can't deal with problems arising out of momentary masking between different parts in a musical arrangement. The side-chaining strategies I explored in Chapter 14 go some way toward tackling the simplest of these situations, but the masking scenarios in musical compositions are typically too complex to be successfully untangled solely using keyed dynamics plug-ins. Masking is a perceptual phenomenon, which means that the human hearing system is essential for detecting it, so it stands to reason that the only truly successful masking compensation comes from manual adjustments evaluated by ear.

These kinds of balance refinements may be subtle, so it makes sense to switch to your Auratone-equivalent while working on them, because that'll cause them to pop out most clearly—they can be very difficult to resolve on stereo nearfields. Chris Lord-Alge also recommends undertaking such tweaks only when listening at low volume: "You can hear better what's going on. When you turn things up, after a while all your moves become a smear. So when you're doing really critical moves, do them at low level."[16] Choose an important instrument and then listen to it for the duration of your mix. If it seems to duck down unduly at any point, edge it up in the mix using automation; if it draws too much attention to itself, then pull it down a fraction. In this kind of application a quarter of a decibel can make a significant difference, although it's also not unusual to see momentary rides of 6dB or more—for example, if

there's a little low-level detail in a note's sustain tail that's disappearing from view.

One of the established masters of this type of automation is hit mixer Andy Wallace, and here's Jim Abbiss describing being a fly on the wall during one of his mixes: "He was constantly adjusting instruments that were playing with the vocals, so you could hear the vocal, and when the vocal stops you can hear the instruments—obvious things which people do, but he did them quite extremely. It meant that he did not need to effect things as much or EQ things as much, because he made space with his balance."[17] The man himself has seldom been interviewed, but on one of those rare occasions, he explained the impetus behind his riding of drum overheads: "I ride them a lot of times because sometimes a cymbal won't be as loud as another cymbals…. But, also… every cymbal crash will be ridden up maybe 5dB or more. Sometimes I will feel that I'm hearing more ambient stuff in the overheads than I want to hear in the mix. [When] the overheads [are] balanced where I want to hear the ambience… sometimes the cymbals won't be loud enough to have the impact that I want."[18]

This is where it really pays off if your faders are set in their area of best control resolution, as I recommended in Section 8.2. If they're right down at the bottom of their travel, you'll struggle to achieve any kind of subtlety with your automation, because you'll only be able to adjust the gain in steps of one or more decibels. That said, personally I often use a gain plug-in for detailed level automation, rather than the channel fader, which circumvents the control-resolution issue. It also makes it easy for me to bypass just that element of the automation for comparison purposes. If you're operating a VST host and fancy giving this a go, then try out GVST's freeware GGain—it provides a ±12dB gain range that is usually perfect for the task.

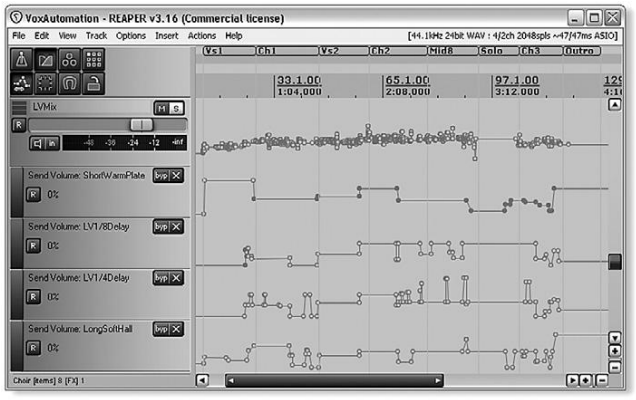

FIGURE 19.6
The vocal level and effect-send automation from a fairly typical Mix Rescue remix. It may look like a lot of data, but that's what it takes to compete in most mainstream musical styles.

It can take a little while to sensitize your ear to subtle balance changes, so don't abandon the listening exercise straight away. Give yourself a few minutes practice to really get into the zone. "The more you listen, the more you hear," says Steve Parr.[19] Nowhere is this more important than with lead vocal parts, of course, so you should listen particularly carefully there. (For more specialist advice on lead vocals, see the "Vocal Rides" box on page 292.) What can sometimes help, particularly when listening on an Auratone substitute, is to think in terms of distance rather than level. If the bass takes a step toward you or away from you from time to time, then it's a clue that rides might be called for.

Directing the Listener's Attention

In any mix, the listener's ear is constantly "refocusing" on different parts as they all jockey for attention. To make a visual analogy, it's a bit like watching a soccer or basketball game on television. While you're focusing on the player in possession, the other players are only in your peripheral vision, and hence are seen in less detail. When the ball's passed to another player, you shift your focus to him instead, and it's then the first player's turn to be in your peripheral vision. The musical listening experience works a bit like this (although hopefully with less dribbling), in that most people will focus most of their attention on only one thing at a time. Your business as the mix engineer is to direct the listener's attention toward what is most important in the track at any given instant. Although the rhythm guitar part might not be nearly as important as the lead vocal, a little fill from that guitar might be much more interesting to listen to than the sound of the singer taking a breath. Fade up just that tiny moment, and it'll demand its own little slice of attention before the vocals return to take pride of place. Gus Dudgeon says, "When you're doing a mix there are slots that appear where you can crank something just enough to help it through so it still makes its point, but isn't blowing your head off and hasn't got lost."[20] "I don't like mixes that are too linear," he adds in another interview, "I like mixes where occasionally something comes out that's a bit of a surprise."[21]

Doing a series of rides like this has several potential benefits. First, the production as a whole will usually appear more involving as a result—in other words, the listener's interest doesn't have as much chance to wane during the lull in vocal activities in our example. "Not every singer has perfect emotion all of the time," says Jon Gass. "So if you can build the mix and craft the textures around the vocal, that's what it's about. It helps sales."[22] Second, the main parts in the mix will be appreciated more deeply because the interludes continually cause you to refocus on them—so you'll concentrate better on the postfill vocal phrase in our example by virtue of coming back to it afresh. These kinds of balance adjustments are, of course, completely unavailable from any kind of mix processor, because they're based on musical judgments rather than on anything a DSP chip can measure, so automation is the only way to bring them about.

My strategy is usually to listen to the song from top to tail, again on my Auratone substitute, and try to think in terms of maintaining the unbroken attention of the listener throughout the timeline, much as Matt Serletic describes here: "A great record should be like you're pulling a string toward you constantly, and you never let up: there's a constant tension that maintains your interest. That can be the vocal dipping into the way the guitar line leads into the prechorus; there's always this constantly connected thread. Especially when I'm doing a final mix, I'm always looking for that point where the string breaks. Then I say 'OK, we've got to fix that.' If I'm not believing it past this point, I've lost it, it's not right."[23]

Don't forget the possibility of riding up reverb or delay tails, because those can be surprisingly arresting in the right circumstances—and they're also something you can add at mixdown if there's nothing else happening in the arrangement!

In a typical song structure, the lead vocal might be responsible for that thread a lot of the time, but wherever this vocal becomes momentarily less interesting (perhaps between phrases, during long note sustains, or after a few repetitions of any simple riff), I'll hunt around among the other arrangement parts to see if there's anything that might provide a welcome diversion. It's in this context that I'll often call on any buried gems I may have turned up during Chapter 5's initial reconnaissance of the multitrack. If they don't appear at a suitable point in the arrangement, I'll consider moving them somewhere else where they're better suited to bridging a gap in the thread of attention. Don't forget the possibility of riding up reverb or delay tails either, because those can be surprisingly arresting in the right circumstances—and they're also something you can add at mixdown if there really is nothing else happening in the arrangement at a given moment!

What's interesting about tinkering with the listener's attention like this is that you can often get away with a surprisingly skewed balance without raising eyebrows, as long as that balance only lasts for a moment and normal service resumes fairly quickly thereafter. This gives you a way to raise the audience's awareness of background parts in your mix, even when they normally have to be balanced quite low. You also have the scope to clearly indicate to the listener that a new instrument has entered the arrangement, even when there's only actually room for the merest suggestion of it in the balance during the remainder of the timeline. If you're relying on the entry of a new part to help differentiate a given section as a way of improving the long-term mix dynamics, then that little dodge can be worth its weight in gold. Another time-honored trick here is to push up the first couple of syllables of a lead vocal part a bit too loud to draw the ear—by the time the listener has had time to think "blimey that vocal's loud!" the vocal has already returned to a less controversial level.

VOCAL RIDES

"With the vocals you're chasing the faders to get them really in your face," says Chris Lord-Alge. "It's all about automation."[24] The amount of time top engineers spend riding lead vocal levels reflects the central importance of singing to most commercial music styles, and it's by no means uncommon for more automation to be applied to lead vocals than to everything else in the track put together—indeed, the data will often look messier than a graph of the Dow Jones. The primary reason is that most listeners remember lyrics and melodies more strongly than anything else, so making sure that both of these elements come through as strongly as possible is fundamental to a track's commercial prospects. "The vocal is probably the most important part of the mix," confirms Tony Visconti. "If the vocal is poorly placed in the mix, it's going to defeat the purpose, it's not going to sell the music. People will remember vocals more than they'll remember the guitar licks…. Don't ever be afraid of putting the vocal too high in the mix…. What you should do is ask someone not involved in the production if they can hear every word. That's really the acid test."[25]

Although maximizing the audibility of lyrics with automation isn't very different from perfecting the balance of any other instrument in a mix, the level of detail required can initially take some getting used to. It's a world where the level of each note, each consonant, each vowel, and each inflection is carefully balanced to create the most solid subjective vocal level possible in the mix. In high-level commercial productions, this activity alone might easily take up several hours, and Bob Clearmountain has even mentioned how über-producer Mutt Lange would spend all day just riding one vocal part.[26] So if you think you're done in ten minutes, think again.

Here are a few pieces of advice on automating for better vocal intelligibility. Dull-sounding consonants often benefit from being faded up quite a bit, because their tone by nature cuts through the mix less. I'm thinking of consonant sounds like "n," "ng," "m," and "l" in particular here, but it depends a bit on the specific singer whether other consonants might also come across unclearly. You can also imply that those consonants are louder than they are by emphasizing the transitions to and from the surrounding vowels. In a similar fashion, consonant sounds "w" and "y" can benefit from an emphasis of their characteristic fast diphthong transition.

But vocal rides aren't just about intelligibility; they're also about maximizing emotional impact, and this is where automating vocals becomes more of an art. All singers have certain vocal characteristics that help to express their emotions, so if you can find and emphasize these moments, you can increase the power of their performance. "It's a great technique," says John Leckie, "because all the little secrets get revealed. At the ends of lines, a lot of singers will trail off, and if you lift the fader 10dB right at the end of the line, there's lots of things you haven't heard before… suddenly there are new things happening in the song."[27] As Leckie suggests, many of the most characterful aspects of vocal performances are to be found in the spaces around the main notes: the little moment just before the vocal properly revs up at the start of a note, the tiny crack in the tone as one pitch changes to another, an expressively extended consonant, a sassy little pitch glide, the unfocused noisy onset of a high-register wail, or the hint of extra breathiness and fall-off as a vulnerable note dies away. Often these kinds of details are hidden within the mix, but if you can unmask them, then you

can significantly elevate the power of the vocal performance—and by association the perceived improvement that your mix has made to the overall production.

The other area where vocal rides can actually appear to improve the vocal performance is in rhythmic styles where the lead vocalist is aiming to deliver a fast line with a good deal of punch and aggression—the most common example these days being the rapping in various up-tempo urban and electronica styles. This kind of performance is tremendously difficult to pull off successfully, especially by an untrained vocalist with underdeveloped diction and breath control. The result is that some syllables may come across well, whereas others will feel rather limp. The solution is to go through and automate a little gain pulse at the onset of each underwhelming syllable. Tedious? Yes, but the alternative is sending your rapper for elocution lessons, which is never going to be pretty.

One final point: vocal automation isn't necessarily just about levels. Many engineers talk about automating de-esser thresholds, for example, while Mike Shipley routinely automates the vocal EQ for his work with Mutt Lange: "With Mutt we always have programmable equalizers where we can EQ every word… every consonant of every word if we want—literally, every part of every word."[28]

Although little fader pushes can really work wonders on a mix, even to the extent of making the music itself appear more engaging, there will occasionally be times when you find you can't make a big enough statement in this way without making a complete mess of the balance. The way to get better results from your rides on these occasions is to simultaneously dip some of the competing background parts. This will decrease masking of the part you're riding up in the mix, while the unchanged levels of more important subsidiary parts in the mix will normally prevent the balance change feeling subjectively awkward. So let's say you really wanted to hear the details of that guitar fill in our vocal/guitar example, but pushing it up too far in the balance felt odd. You could draw more attention to it by pulling back the faders of the other background parts in the mix at that moment (say piano, synth pad, and backing vocal lines) to clear more space for the fill's qualities to shine through at a lower level. The beauty of doing it like this is that as long as the vocal, drums, and bass remain at their usual levels, the listeners probably won't even be conscious of what you've done at all, because their attention is targeted at the vocal and guitar lines. In some such cases, you may find that you can completely mute some background parts without anyone being the wiser—all people will notice is that the instruments they're supposed to be listening to sound clearer. Result!

19.6 NAILING DOWN THE FINAL VERSION

Once your automation work is done, you'll have arrived at a complete mix draft. Only two things now stand between this and the final mixdown file: another

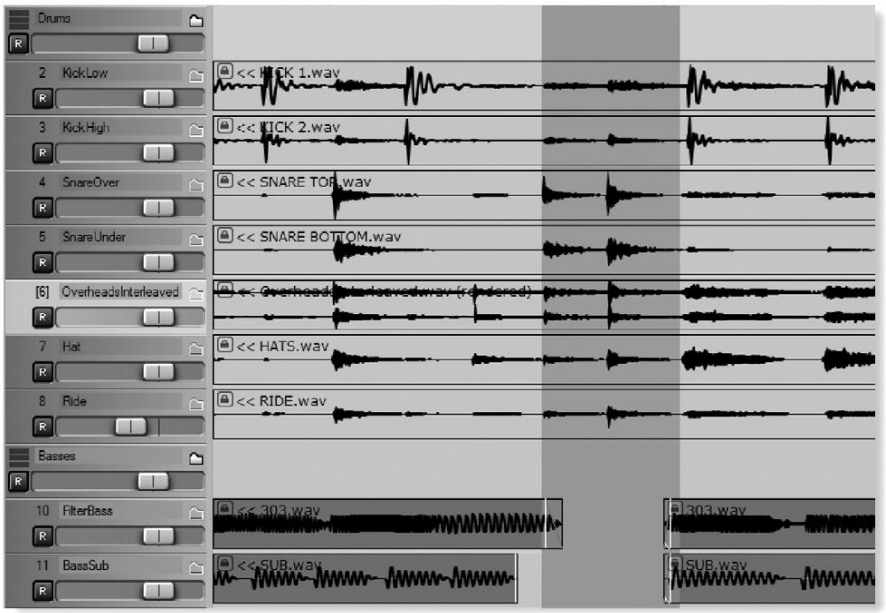

FIGURE 19.7
Clearing out uninteresting backing moments behind important foreground musical features can help the latter come through more clearly. In this screenshot, for example, I've edited out the fuzzy synth bass under an important snare fill.

bout of mix referencing, and the process of incorporating any revisions suggested by the artist/client. (If it's your music, of course, then both these tasks may be rolled into one.) As far as the referencing stage goes, you should first of all proceed with the same kinds of A/B comparisons you performed with the first mix section in Section 19.3. However, there's one further referencing activity I'd recommend, which really helps to nail down the final subtle tweaks: building a "snag list." This involves listening to the whole mix on all your different monitors and making notes of any final nips and tucks required. Although that sounds similar to the mix referencing you've already done, it actually works a bit differently, and the devil's in the detail.

Effective Snagging

You should now make a point of listening to your mix (and any reference tracks) from top to tail, in order to cement your decisions regarding the long-term mix dynamics—a lot of your mixing time will inevitably have been spent playing back small snippets, so a good few complete play-throughs at this stage acts as a valuable counterweight. "My usual rule," says Andy Bradfield, "is that if I play a track three times in a row and something consistently bugs me, then I'm going to have to do something about it. That's how I know that something isn't right—my ear starts complaining and my brain keeps bringing it up."[29] Bob Clearmountain imparts some further words of wisdom: "Make sure you're

listening to the whole mix, not just the vocal, the guitar, or the snare drum…. You have to train yourself to listen to the song and not the individual parts, and then things that are wrong will jump out at you—it'll be obvious that there's a part in the second verse that's pulling your ear away from the vocal when the lyric at that point is really important to the song."[30]

There's no better snagging time than first thing in the morning after a good night's sleep. Keeping your concentration levels up throughout even a single complete mix playback is demanding, so it's vital that your ears and brain are as fresh as possible before you start. "The hardest thing is trying to stay fresh, to stay objective," says Nigel Godrich. "It's so hard to let it all go and become a listener"[31] I'd also encourage you not to play the mix from your DAW system, either, but instead to burn a CD (or export the mixdown file to your portable music player) so that you can listen to it in exactly the same way you listen to your own record collection. This helps you trick yourself psychologically into listening more as a consumer than as a mix engineer, which is a remarkably powerful way of highlighting when a mix is truly fit for general consumption. Within this context, a couple of playbacks on your favorite day-to-day domestic listening systems can be a good idea. Although your hi-fi, car stereo, or iPod headphones won't actually provide you with much extra information about the technicalities of your mix if you've used your monitoring systems wisely thus far, they can make it easier to appreciate your production from Joe Public's perspective.

An additional benefit of playing from a CD or bounced file is that you're less likely to be distracted by visual feedback from your DAW's meters, waveform displays, and arrangement-window layout. "It's real easy to be objective when you're not staring at the speakers and looking at the meters," says Ed Seay.[32] If you can see your software's play cursor advancing to a busier-looking section of the onscreen arrangement labeled "chorus," you'll instinctively be inclined to think that this section sounds bigger when it arrives, even if it's actually falling a bit flat. Furthermore, because you're preventing yourself tweaking the DAW's mix settings while listening, you'll typically be much more disciplined in basing your decisions on the consensus of all your different monitoring systems—it's all too easy to start fiddling with mix settings simply to compensate for the vagaries of a particular monitoring system, when there's nothing actually wrong with the mix. Gareth Jones recalls making this mistake with a Depeche Mode mix early in his career: "Since the record company was so concerned about getting the song on the radio we spent a lot of time listening and mixing on a two- or three-inch speaker…. We probably spent too much time doing this, because while it sounded absolutely amazing on a small system, on bigger systems I can hear all the faults. You know, 'Oh God, I could have got more bass on it,' or 'It would have been great if the

FIGURE 19.8
For the final snagging stage of your mix evaluation, burn your mix to CD. That way you'll tend to listen to it more as a consumer than as an engineer, which is a useful perspective.

high end could have been a bit smoother.' Obvious things. If I'd been a more experienced engineer, I would have… also paid a bit more attention to the big monitors."[33] In my experience, there's little danger of falling into this kind of trap if you make a point of whittling down your list of snags on paper first, leaving only those that are truly justifiable across all your monitoring systems.

A snag list therefore helps you wrap up a mix much more efficiently and decisively, because you always have a definite target in your sights while working and it's very difficult for niggly mix mistakes to slip through the net. When you've ticked off all the snags on your list, and referenced again to check they've been properly dealt with, then you'll be much more confident in consigning your mix to the out tray. If you still worry that your mix is unfinished after all that, then perhaps it's time to just resign yourself to it. Phil Tan notes, "If you were to give me the same song on different days, it would probably come out differently. There's no such thing as perfection, and I think music is more fun when things are a little inexact."[34]

Revision Requests

Of course, if you're mixing someone else's music, then it's neither here nor there whether you're personally happy with the mix, because it's the client who decides what's finished and what isn't. No matter how good an engineer you are, there's simply no getting away from the fact that a lot of your mix decisions will effectively be guesswork—calculated guesses, admittedly, but guesses nonetheless. Without being Marvo the Incredible Mind Reader, you can't know exactly how the artist, manager, or A&R rep want the record to sound, even if you've taken the wise precaution of asking them for some representative reference tracks for guidance. Therefore, even at the highest levels of the industry, mix engineers expect (and plan for) the eventuality that they may need to revise their first submitted mix in response to suggestions from the client. This process can be inconvenient or vexing at times, but it's inevitable, so there's no sense in getting upset about it (grrr!), even if you find you have to go right back to the drawing board.

Mixing isn't just a technical exercise; it's a process of discovering what the client's looking for, and sometimes you need to try a couple of blind alleys before you can find exactly the right direction. Marcella Araica, like most high-profile mix engineers, has been there and bought the T-shirt: "When you're recording, people are usually satisfied just hearing the sound come back from the monitor speakers, but when it comes to the mix it's suddenly the real deal… you get a lot of cooks in the kitchen, which can be frustrating. But it's part of the whole mix thing. I find that I can learn a lot when people have comments. They may be hearing something that I don't, and even when the critique makes no sense, it may lead me to bring something out that I hadn't heard before."[35] Daniel Lanois picks up on the same theme: "Just go with the other person's idea. The thing about ideas is, if you don't chase them up and see them through to their

conclusion, you may have frustration living in a corner of the room. And I've never had a disagreement regarding result. Somehow it always works out that everybody agrees on the approach that sounds best."[36]

Making the process of revisions as painless as possible is as much psychology as anything, so it will always depend heavily on the kinds of people you're dealing with. However, Kyle Lehning takes an approach that has proved its worth for me on many occasions: "I tell the artist very early on in the process that there will never be a tie. 'Whatever we find ourselves disagreeing on, you win!' But because there won't be a tie, I have the freedom to tell you exactly what I think. At the end of the day, it's your career, so it will be your responsibility to decide whether what I have to say is helpful or not…. Sometimes people are afraid of losing their job or offending the artist, but the more successful an artist gets, the more important it is for them to find somebody who will actually tell them what they really think."[37]

Beyond that general advice, though, I do have a few further suggestions that have helped me navigate this particular minefield in the past. The main one is to ask your clients to describe the changes they're after with relation to other productions (as suggested at the end of Section 4.1). So if the vocal reverb isn't right, ask them to play you something more suitable from their record collection. This heads off two frequent problems: the first is that opinions may differ regarding what a "silky vocal reverb" is, and the second is that what they call "reverb" might actually be delay, chorusing, double-tracking, or any number of other things. To misquote a well-known adage: an audio example is worth a thousand words! It can also be useful to refer back to any rough mix that the client supplied and find out how close that sound was to what they were looking for.

Ask your clients to describe the changes they're after with relation to other productions. So if the vocal reverb isn't right, ask them to play you something more suitable from their record collection. An audio example is worth a thousand words!

"I can't assume that my sensibility is the right sensibility," remarks Justin Niebank. "That's where going back and listening to rough mixes can be especially useful—they can tell you what people were thinking."[38] Jaycen Joshua is more blunt: "[They've] spent so much time and so much creative energy on these rough mixes that if you don't understand the track and their ideas, you are going to get killed."[39] In some cases, you can find that a specific effect is vital to the client's vision of the track, in which case the quickest results are often to be had simply by bouncing that effect from the rough mix or finding out the exact settings so that you can go about replicating it as closely as possible.

My final recommendation when dealing with client revisions is to keep backups of the project and mixdown files for each stage of revisions, because it's not uncommon for people to backtrack on a revision idea once they actually

hear what it does to the mix in practice. And it should go without saying that keeping backups of all your work in general is essential. My maxim is this: files do not really exist unless they're saved in two different places at once, and they're only properly backed up when I can't stupidly delete them by mistake!

CUT TO THE CHASE

- If you're working in a style of music that relies on the effects of buss compression, then you should get that processing involved from the moment all the tracks and effects in your first mix section are happening, because you may need to make balance compensations in response to the gain-reduction action. Resist the temptation to use frequency-selective dynamics processing in this role—stick to full-band processors and concentrate on minimizing unwanted side effects.
- If you chicken out of referencing your mix properly, then you'll never learn how to make your mixes compete with commercial records. Be prepared to go through several stages of referencing and readjustments. Make sure to compare the mix's overall tonality and stereo spread; the balance, tone, and stereo image of the most important instruments; and the levels of any reverb and delay effects.
- In order to reference your mix against aggressively loudness-enhanced commercial material, you may need to apply loudness processing for comparison processes. There are lots of methods of loudness maximization, and a combination of these techniques often gives the best results. When evaluating the suitability of the processing, remove the loudness boost from the equation and listen carefully for processing side effects using all your monitoring systems.
- Building up the complete mix from your first completed mix section demands no new processing or effects techniques, but it will usually require dynamic mixer automation to maintain the best balance throughout the production. If you need to create any dramatic buildups, remember that you can only rev things up if you first drop them down. Detailed level rides are essential for most productions, because they can direct the listener's attention where you want it and can also ensure the maximum audibility of lead vocals and important hooks.
- As a last stage of referencing, build up a written snag list, because this helps focus your mind on what is necessary for completion. For best results, do this in the morning following a night's rest, listen to the track all the way through, avoid looking at visual displays while you're listening, and evaluate the validity of any snags on all your different monitoring systems (including your favorite domestic playback systems) before actually implementing any changes.
- Even for the top professionals, mix revisions are a fact of life if you work on other people's music, so don't expect the client to approve your first-draft mix straightaway. If possible, try to communicate about the mix in terms of concrete audio examples from commercial records, and make sure to keep backups of every draft as you go in case anyone decides to backtrack on an idea.

Assignment

- Apply buss compression to your mix if it's appropriate, and readjust the balance after that if necessary. Then export that section of the mixdown as a stereo audio file and import it to a new DAW project for referencing against relevant commercial tracks. If loudness differences are a problem, then consider using loudness-enhancement processing on the mix at this stage for referencing purposes. If the overall tonality of the mix is wrong, then try to correct it with a single high-quality EQ plug-in over the main mix buss in your mixdown project, rather than messing with too many of your individual channel processors.
- When you've completed the referencing process and revised your first mix snippet, build up the rest of the arrangement sections in order of importance, engaging automation where necessary to achieve a full mix draft. Export an audio file of the full mix draft and reference against commercial productions again. To finish off, leave the mix overnight and build a snag list the following morning. Once those snags are dealt with, seek the client's feedback and implement any suggestions. When the mixdown's finally a wrap, generate any alternate mix versions and make sure you've backed everything up.

http://www.cambridge-mt.com/ms-ch19.htm

This page intentionally left blank

CHAPTER 20

Conclusion

In this book I've tried to cover the main areas of mix technique that small-studio users need to master to achieve commercial-level results. Essentially it all boils down to the following:

- Setting up and learning to use your monitoring systems
- Organizing, comping, arranging, and correcting the raw recordings
- Balancing all the tracks and adjusting their timbres if necessary
- Adding sweetening effects, buss compression, and automation
- Referencing your work, and creating any necessary alternate mix versions

I've deliberately taken a step-by-step approach in order to make sense of so many different subjects and also so that you can build up your own mix while working through the book. However, as I mentioned in Chapter 15, most real-world mix engineers don't actually follow a linear trajectory and instead react much more fluidly to the mix they're hearing from moment to moment. So while a structured approach can be reassuring at first, don't be afraid to deviate from that model as your confidence increases, because that's the key to developing your own unique mixing style.

I've covered a lot of ground in this book, and if you've reached this point you should already have improved your mixing by learning how to apply some of the top pro techniques within your own small studio. The only observation I still want to add is this: learning to mix takes time and effort. I've done what I can to speed up the process for you, but you'll still need to put in some hard graft if you're going to hold your own among the top names in the business.

"I'm all for people getting into recording music, because it's a great thing to do," says Tony Platt, "but if that is your chosen career path then you should spend a bit of time and work out how to do it properly. When I say properly I'm not standing in judgment saying that's right or that's wrong… but it should be approached as a skill and you should take it as far as you possibly can. Don't approach it with the attitude of 'Oh well, that'll do,' because that's not good enough."[1] Tony Maserati stresses the point: "I studied very, very hard. I would be on the subway with my SSL manual, reading it and marking it. I tell my assistants now: it's great to record your

Mixing Secrets for the Small Studio.

friends' bands, but don't just do that. Take a tape or a file and the manuals and work through a room's gear—the reverbs, the plug-ins. Keep working it over and over. I'd sit there with the gear for hours. I think that helped me become a mixer, and to be able to come up with things that were creative and new."[2]

For some final words of encouragement, let me leave you with Al Schmitt: "Stick with what you're doing, hang in there, and keep doing it…. I can't tell you how many times guys get turned down and rejected and battered about, and then all of a sudden, wham!—they just pop through. I don't know a better way to enjoy your life than making a living doing something you really love to do."[3]

APPENDIX 1

Who's Who: Selected Discography

Here's a list of all the engineers I've cited in this book, along with some of the most influential recordings they've been involved with.

ABBISS, JIM

Arctic Monkeys: *Whatever People Say I Am, That's What I'm Not*; Ladytron: *Witching Hour*; Sneaker Pimps: *Becoming X*; Editors: *The Back Room*; Kasabian: *Kasabian, Empire*; Adele: *19*; KT Tunstall: *Tiger Suit*.

AINLAY, CHUCK

Trisha Yearwood: *Where Your Road Leads*, *Thinkin' about You*, *Everybody Knows, How Do I Live, Real Live Woman*; George Strait: *Somewhere Down in Texas*, *Pure Country*, *Blue Clear Sky*, *Carrying Your Love with Me*; Vince Gill: *High Lonesome Sound, The Key*; Dixie Chicks: *Wide Open Spaces*; Mark Knopfler: *Sailing to Philadelphia, The Ragpicker's Dream, Wag the Dog, Metroland*; Patty Loveless: *If My Heart Had Windows, Honky Tonk Angel*.

ALBINI, STEVE

The Pixies: *Surfer Rosa*; Nirvana: *In Utero*; Bush: *Razorblade Suitcase*; PJ Harvey: *Rid of Me*; Jimmy Page and Robert Plant: *Walking into Clarksdale*; Manic Street Preachers: *Journal for Plague Lovers*; The Wedding Present: *El Rey*; The Breeders: *Mountain Battles*; The Stooges: *The Weirdness*; Cheap Trick: *Rockford*.

ARAICA, MARCELLA

Britney Spears: *Blackout*, *Circus*; Timbaland: *Shock Value*; Mariah Carey: $E = MC^2$; Danity Kane: *Welcome to the Doll's House*; Keri Hilson: *In a Perfect World*; Black Eyed Peas: *Monkey Business*.

AVRON, NEAL

Fall Out Boy: *From under the Cork Trees*, *Infinity on High*, *Folie a Deux*; Weezer: *Make Believe*; Everclear: *So Much for the Afterglow*; Linkin Park: *Minutes to Midnight*; The Wallflowers: *Bringing Down the Horse*; Lifehouse: *No Name Face*; Yellowcard: *Ocean Avenue*.

BALLARD, GLEN

Wilson Phillips: *Wilson Phillips*; Alanis Morissette: *Jagged Little Pill, Supposed Former Infatuation Junkie*; Michael Jackson: *Thriller, Bad, Dangerous*; POD: *Testify*; Anastacia: *Anastacia*; Goo Goo Dolls: *Let Love In*; Christina Aguilera: *Stripped*; No Doubt: *Return of Saturn*; Aerosmith: *Nine Lives*; Paula Abdul: *Forever Your Girl*.

BARRESI, JOE

Queens of the Stone Age: *Queens of the Stone Age*, *Lullabies to Paralyze*; Tool: *10,000 Days*; The Melvins: *Stag, Honky*; Hole: *Celebrity Skin*; Limp Bizkit: *Chocolate Starfish & the Hotdog Flavored Water*; The Lost Prophets: *Start Something*; Skunk Anansie: *Stoosh*; Bad Religion: *The Empire Strikes First, New Maps of Hell, The Dissent of Man*.

BLAKE, TCHAD

Sheryl Crow: *The Globe Sessions*; Bonnie Raitt: *Souls Alike*; Phish: *Undermind*; Suzanne Vega: *Nine Objects of Desire, 99.9°F, Beauty and Crime*; Crowded House: *Farewell to the World*; Paul Simon: *Surprise*; Peter Gabriel: *Ovo, Long Walk Home, Up*; The Dandy Warhols: *Come Down, Odditorium*; Neil Finn: *One Nil, Try Whistling This*; Pearl Jam: *Binaural*.

BOTTRELL, BILL

The Traveling Wilburys: *The Traveling Wilburys Volume 1*; Tom Petty: *Full Moon Fever*; Madonna: *Like a Prayer*; The Jacksons: *Victory*; Michael Jackson: *Bad, Dangerous*; Sheryl Crow: *Tuesday Night Music Club*, *Sheryl Crow*; Shelby Lynne: *I Am Shelby Lynne*; Elton John: *Songs from the West Coast*.

BRADFIELD, ANDY

Robbie Robertson: *Contact from the Underworld of Redboy*; Rufus Wainwright: *Want One*, *Want Two*, *Release the Stars*; Josh Groban: *Awake*; Alanis Morissette: *Flavors of Entanglement*; Spice Girls: *Spice*.

BRAUER, MICHAEL

Coldplay: *Parachutes, X&Y, Viva La Vida*; Athlete: *Tourist*; The Fray: *The Fray*; Evans Blue: *The Pursuit Begins When This Portrayal of Life Ends*; Paolo Nutini: *Last Request*; John Mayer: *Continuum*; James Morrison: *Undiscovered*; The Kooks: *Inside In/Inside Out*; My Morning Jacket: *Evil Urges*; Travis: *The Boy with No Name*; The Doves: *Kingdom of Rust*; KT Tunstall: *Drastic Fantastic*; Idlewild: *Warnings/Promises*; Fountains of Wayne: *Traffic and Weather*; Aimee Mann: *Lost in Space*; Bob Dylan: *Lovesick*.

BUSH, STEVE

Stereophonics: *Word Gets Around, Performance & Cocktails, Just Enough Education to Perform*; Corinne Bailey Rae: *Corinne Bailey Rae*.

CAILLAT, KEN

Fleetwood Mac: *Rumours, Tusk, Mirage*; Colbie Caillat: *Coco, Breakthrough*; Christine McVie: *In the Meantime*; Lionel Ritchie: *Dancing on the Ceiling*; Michael Jackson: *Bad*.

CASTELLON, DEMACIO

Madonna: *Hard Candy, Celebration*; Nelly Furtado: *Loose, Mi Plan*; Rihanna: *Good Girl Gone Bad*; Monica: *After the Storm*; The Game: *The Documentary*.

CHERNY, ED

The Rolling Stones: *Stripped*, *Bridges to Babylon*, *No Security*; Bonnie Raitt: *Nick of Time*, *Luck of the Draw*, *Longing in Their Hearts*, *Road Tested*; Bob Dylan: *Under the Red Sky*, *Unplugged*.

CHICCARELLI, JOE

White Stripes: *Icky Thump*; Frank Zappa: *Sheik Yerbouti, Joe's Garage, Tinseltown Rebellion*; My Morning Jacket: *Evil Urges*; The Shins: *Wincing the Night Away*; Counting Crows: *The Desert Life*.

CHURCHYARD, STEVE

The Pretenders: *Learning to Crawl*; Counting Crows: *Recovering the Satellites*; Celine Dion: *Falling into You*; Ricky Martin: *Vuelve*, *Almas Del Silencio*, *Sound Loaded*; Shakira: *Laundry Service*; The Stranglers: *La Folie*, *Feline*; Big Country: *Wonderland*; Bryan Ferry: *Boys & Girls*; INXS: *Listen Like Thieves;* Kelly Clarkson: *Thankful*.

CLEARMOUNTAIN, BOB

The Pretenders: *Get Close*; Bryan Adams: *Into the Fire*, *Reckless*, *Cuts Like a Knife, So Far So Good, 18 till I Die, Room Service, 11*; Bruce Springsteen: *Born in the USA*; The Rolling Stones: *Tattoo You*; Bon Jovi: *These Days, Bounce, Crush*; Roxy Music: *Avalon*; David Bowie: *Let's Dance*; INXS: *Kick, Full Moon, Dirty Hearts, Welcome to Wherever You Are;* The Corrs: *Talk on Corners, Forgiven Not Forgotten, Unplugged*; Robbie Williams: *Intensive Care*; Simple Minds: *Once Upon a Time, Black & White 050505, Graffiti Soul*; Sheryl Crow: *Wildflower*; Aimee Mann: *Whatever, Bachelor #2*; Rufus Wainwright: *Rufus Wainwright*.

CLINK, MIKE

Guns 'n' Roses: *Appetite for Destruction*, *GNR Lies*, *Use Your Illusion*, *The Spaghetti Incident*; Survivor: *Eye of the Tiger*; Megadeth: *Rust in Peace*; Mötley Crüe: *New Tattoo*; Whitesnake: *Whitesnake.*

COSTEY, RICH

Muse: *Absolution, Black Holes & Revelations*; Interpol: *Our Love to Admire*; Polyphonic Spree: *Together We're Heavy*; Franz Ferdinand: *You Could Have It So Much Better*; The Mars Volta: *Frances the Mute, Deloused in the Comatorium*; Audioslave: *Audioslave*; Weezer: *Make Believe*; Bloc Party: *Silent Alarm*; Doves: *Some Cities*; My Chemical Romance: *Three Cheers for Sweet Revenge*; Three Days Grace: *Three Days Grace*; Jimmy Eat World: *Futures*; POD: *Payable on Death*; Rage against the Machine: *Renegades*; Fiona Apple: *When the Pawn*; Jurassic 5: *Quality Control*; Nine Inch Nails: *With Teeth*.

DAVIS, KEVIN "KD"

Coolio: "Gangsta's Paradise"; Destiny's Child: *The Writing's on the Wall*, *Survivor*; N'Sync: *No Strings Attached*; Pink: *Missundaztood*; Tupac Shakur: *Me against the World*; TLC: *3D*; Usher: *8701*, Outkast: *Speakerboxx/The Love Below.*

DE VRIES, MARIUS

Massive Attack: *Protection*; Bjork: *Debut, Post, Homogenic, Vespertine*; Madonna: *Bedtime Stories, Ray of Light*; U2: *Pop*; Mel C: *Northern Star*; David Gray: *White Ladder*; Rufus Wainwright: *Want One, Want Two*; Josh Groban: *Awake.*

DECARLO, LEE

John Lennon: *Double Fantasy*; Black Sabbath: *Live Evil*; Quiet Riot: *Quiet Riot II.*

DORFSMAN, NEIL

Dire Straits: *Brothers in Arms, Love over Gold*; Bruce Hornsby: *Scenes from the South Side*; Paul McCartney: *Flowers in the Dirt*; Bruce Springsteen: *Tracks*, *The River*; Sting: *Nothing Like the Sun*, *Brand New Day*; Bjork: *Medulla.*

DOUGLAS, JACK

Aerosmith: *Get Your Wings*, *Toys in the Attic*, *Rocks*, *Draw the Line*, *Rock in a Hard Place*, *Honkin' on Bobo*; John Lennon: *Imagine, Double Fantasy*; Cheap Trick: *Cheap Trick*, *At Budokan*, *Standing on the Edge.*

DOUGLASS, JIMMY

Aaliyah: *One in a Million*, *Are You That Somebody*; Missy Elliott: *Supa Dupa Fly*, *Da Real World*, *Miss E… So Addictive* (including "Get Ur Freak On"); Aretha

Franklin: *Young, Gifted & Black*; Jay-Z: *Hard Knock Life*; Jodeci: *The Show, The Hotel, The After Party*; NERD: *In Search Of*; Timbaland & Magoo: *Timbaland and Magoo, Tim's Bio, Up Jumps the Boogie, Love 2 Love U, All Y'All.*

DRESDOW, DYLAN "3D"

Black Eyed Peas: *Elephunk, Monkey Business, The END*; Flo Rida: *ROOTS*; The Game: *Doctor's Advocate*; Common: *Finding Forever*; Christina Aguilera: *Stripped*; Methodman & Redman: *Blackout, How High*.

DUDGEON, GUS

Elton John: *Elton John, Madman across the Water, Honky Chateau, Goodbye Yellow Brick Road, Caribou, Captain Fantastic and the Brown Dirt Cowboy*; David Bowie: *Space Oddity*; Original Soundtrack: *Tommy*.

ELEVADO, RUSS

D'Angelo: *Voodoo*; Al Green: *Lay It Down*; Alicia Keys: *Songs in A Minor, Diary of Alicia Keys*; The Roots: *Phrenology, Game Theory*; Dru Hill: *Dru Hill*: Erykah Badu: *Mama's Gun*; Boney James: *Ride*; Blackalicious: *Blazing Arrow, The Craft*.

ELMHIRST, TOM

Amy Winehouse: *Frank, Back to Black*; Adele: *19*; Goldfrapp: *Black Cherry*; Lily Allen: *Alright, Still*; Paolo Nutini: *These Streets*; Mark Ronson: *Version*; Sugababes: *Three*.

EMERICK, GEOFF

The Beatles: *Revolver, Sgt. Pepper's Lonely Hearts Club Band, Magical Mystery Tour, Abbey Road*; Paul McCartney: *Band on the Run, Flaming Pie*; Elvis Costello: *Imperial Bedroom, All This Useless Beauty*; Badfinger: *No Dice*; Robin Trower: *Bridge of Sighs.*

ENDERT, MARK

Maroon 5: *Songs about Jane, It Won't Be Soon before Long*; Miley Cyrus: *Breakout*; Delta Goodrem: *Delta*; Secondhand Serenade: *Hear Me Now*; Rihanna: *Good Girl Gone Bad*; Madonna: *Ray of Light, Music*; Quietdrive: *When All That's Left Is You*; Tracey Bonham: *Down Here*; Fiona Apple: *Tidal*; Gavin DeGraw: *Chariot.*

FILIPETTI, FRANK

Foreigner: *Agent Provocateur, Inside Information*; Hole: *Celebrity Skin*; Korn: *Untouchables, Here to Stay*; Carly Simon: *Hello Big Man, Spoiled Girl, Coming around Again, Film Noir, Bedroom Tapes*; Barbra Streisand: *Higher Ground*; James Taylor: *Hourglass, That's Why I'm Here*; Kiss: *Lick It Up*; Bangles: *Everything*; Survivor: *Too Hot to Sleep.*

FINN, JERRY

Alkaline Trio: *Crimson, Remains, Good Mourning*; Green Day: *Kerplunk, Dookie, Insomniac*; Smoking Popes: *Destination Failure*; MxPx: *The Ever Passing Moment, Ten Years and Running*; Blink-182: *Enema of the State, Take off Your Pants and Jacket, Blink-182*; Tiger Army: *Music from Regions Beyond*; +44: *When Your Heart Stops Beating*; AFI: *Sing the Sorrow, Decemberunderground, Crash Love*; Sparta: *Wiretap Scars*.

FLOOD

Editors: *In This Light and on This Evening*; 30 Seconds to Mars: *This Is War*; Goldfrapp: *Seventh Tree*; PJ Harvey: *White Chalk, Is This Desire, To Bring You My Love*; The Killers: *Sam's Town*; U2: *The Joshua Tree, Achtung Baby, Pop, Zooropa, How to Dismantle an Atomic Bomb*; Smashing Pumpkins: *Mellon Collie and the Infinite Sadness, Machina/Machines of God, Adore*; Nine Inch Nails: *Downward Spiral, Pretty Hate Machine*; Depeche Mode: *Violator, Pop Will Eat Itself, Songs of Faith and Devotion*; Nick Cave and the Bad Seeds: *From Her to Eternity, The Firstborn Is Dead, Kicking against the Pricks, Your Funeral My Trial, Tender Prey, The Good Son*; Erasure: *Wonderland, The Circus.*

GABRIEL, PASCAL

Goldfrapp: *Head First*; Marina and the Diamonds: *The Family Jewels*; Ladyhawke: *Ladyhawke*; Miss Kittin: *BatBox*; Little Boots: *Hands*; Dido: *No Angel*; Kylie Minogue: *Fever*.

GASS, JON

Paula Abdul: *Forever Your Girl*; Madonna: *Bedtime Stories*; Dru Hill: *Enter the Dru*; Usher: *My Way*; Whitney Houston: *Just Whitney*; Michael Jackson: *Invincible*; Mary J. Blige: *Mary*; Destiny's Child: *The Writing's on the Wall.*

GATICA, HUMBERTO

Michael Jackson: *Bad*; Celine Dion: *Celine Dion, D'eux, Falling into You, Let's Talk about Love, A New Day Has Come, One Heart*; Michael Bublé: *Call Me Irresponsible, Crazy Love*; Ricky Martin: *Vuelve*; Cher: *Believe*; Julio Iglesias: *Non Stop, Un Hombre Solo, Crazy, La Carretera, Tango*; Barbra Streisand: *The Mirror Has Two Faces, Higher Ground.*

GHENEA, SERBAN

Katy Perry: *One of the Boys, Teenage Dream*; Ke$ha: *Animal*; Jason Derulo: *Jason Derulo*; Miley Cyrus: *The Time of Our Lives*; Britney Spears: *In the Zone, Circus*; Iyaz: *Replay*; Jay Sean: *All or Nothing*; Weezer: *Raditude*; The Fray: *The Fray*; Kelly Clarkson: *All I Ever Wanted, Breakaway*; Pink: *I'm Not Dead, Funhouse*; Avril

Lavigne: *The Best Damn Thing*; Black Eyed Peas: *Monkey Business*; Fergie: *The Dutchess*; Carrie Underwood: *Some Hearts*; Justin Timberlake: *Justified*; Mark Ronson: *Version*; Sugababes: *Change*; Jill Scott: *Who Is Jill Scott? Beautifully Human, The Real Thing*; Dave Matthews Band: *Stand Up*; Robbie Williams: *Rudebox*; Usher: *My Way*; NERD: *In Search Of*.

GLOSSOP, MICK

Frank Zappa: *Shut up N Play Yer Guitar, Joe's Garage*; John Lee Hooker: *Don't Look Back*; Van Morrison: *Too Long in Exile, Hymns of the Silence, Enlightenment, Poetic Champions Compose, No Guru No Method No Teacher, Inarticulate Speech of the Heart, Into the Music, Wavelength*; The Waterboys: *This Is the Sea, The Whole of the Moon*; Tangerine Dream: *Rubycon, Ricochet*; The Wonder Stuff: *Never Loved Elvis, The Size of a Cow, Welcome to the Cheap Seats, Dizzy, Cursed with Insincerity*; Lloyd Cole: *Antidepressant, Music in a Foreign Language, Broken Record*.

GODRICH, NIGEL

Radiohead: *The Bends, OK Computer, Kid A, Amnesiac, Hail to the Thief, In Rainbows*; Thom Yorke: *The Eraser*; Natalie Imbruglia: *Left of the Middle*; Beck: *Mutations, Sea Change, Guero, The Information*; The Divine Comedy: *Regeneration, Absent Friends*; Travis: *The Man Who, The Invisible Band, The Boy with No Name*; Paul McCartney: *Chaos and Creation in the Backyard*; Air: *Talkie Walkie, Pocket Symphony*; Pavement: *Terror Twilight*; REM: *Up*.

GOLDSTEIN, JASON

Kelly Rowland: *Ms Kelly*; Rihanna: *Music of the Sun*; Ludacris: *Theatre of the Mind*; The Roots: *Game Theory, Rising Down*; Jay-Z: *The Blueprint*; R Kelly & Jay-Z: *Unfinished Business*; Nas: *Stillmatic*; Mary J. Blige: *Mary*; Beyoncé: *B'Day*.

GUZAUSKI, MICK

Eric Clapton: *Back Home*; Christina Aguilera: *Christina Aguilera*; Babyface: *The Day*; Backstreet Boys: *Millennium*; Boyz II Men: *II*; Toni Braxton: *Secrets* (including "Unbreak My Heart"); Brandy: *Never Say Never*; Mariah Carey: *Music Box* (including "Dream Lover" and "Hero"), *Butterfly, Charmbracelet*; Patti LaBelle: *Winner in You*; Jennifer Lopez: *On the 6, J.Lo, This Is Me… Then*; Barbra Streisand: *Higher Ground*; Leann Rimes: *You Light up My Life*.

HODGE, STEVE

Janet Jackson: *Control, Rhythm Nation, Janet, The Velvet Rope, All for You, Damita Jo*; George Michael: *Faith*; Mariah Carey: *Rainbow, Share My World, Glitter, Butterfly, Charmbracelet*; Michael Jackson: *HIStory*; Toni Braxton: *More Than a Woman*.

JEAN, WYCLEF

The Fugees: *Blunted on Reality, The Score*; Wyclef Jean: *The Carnival, The Ecleftic, Masquerade, The Preacher's Son*; Santana: *Supernatural*; Destiny's Child: *Destiny's Child*; Pras: *Ghetto Supastar, Win Lose or Draw*; Mya: *Fear of Flying*; Michael Jackson: *Blood on the Dance Floor*; Shakira: *Oral Fixation Vol. 2; She Wolf.*

JOHNS, ANDY

Led Zeppelin: *Led Zeppelin III*, *Led Zeppelin IV*, *Houses of the Holy*, *Physical Graffiti*; The Rolling Stones: *Sticky Fingers*, *Exile on Main Street*, *Goat's Head Soup*; Van Halen: *For Unlawful Carnal Knowledge*; Television: *Marquee Moon*; Jethro Tull: *Stand Up.*

JONES, GARETH

Erasure: *Wild!*, *Erasure, Cowboy, Other Peoples Songs, Light at the End of the World*; Depeche Mode: *Construction Time Again, Some Great Reward, Exciter*; Wire: *The Ideal Copy, A Bell Is a Cup… Until It Is Struck.*

JONES, QUINCY

Frank Sinatra: *L.A. Is My Lady, It Might as Well Be Swing*; Count Basie: *This Time by Basie*; Michael Jackson: *Off the Wall, Thriller, Bad*; Quincy Jones: *Walking in Space, Smackwater Jack, The Dude, Back on the Block, Q's Jook Joint, You've Got It Bad Girl*; George Benson: *Give Me the Night*; USA for Africa: "We Are the World."

JOSHUA, JAYCEN

Mariah Carey: *Memoirs of an Imperfect Angel, E = MC*^{2}; Beyoncé: *I Am Sasha Fierce*; Jamie Foxx: *Intuition*; Sean Paul: *The Trinity*; Justin Bieber: *My World 2.0*; Rihanna: *Rated R*; Mary J. Blige: *Growing Pains, Stronger with Each Tear*; Christina Aguilera: *Bionic*; New Kids on the Block: *The Block*; Soulja Boy: *Souljaboytellem.com*; Pussy Cat Dolls: *Doll Domination*; Usher: *Here I Stand*; Monica: *Still Standing*; Ashanti: *The Declaration.*

KATZ, BOB

Paquito D'Rivera: *Portraits of Cuba*; Olga Tañon: *Olga Viva, Viva Olga*; Sinead O'Connor: *Theology.*

KILLEN, KEVIN

Peter Gabriel: *So*; U2: *War*, *The Unforgettable Fire*, *Rattle and Hum*; Elvis Costello: *Spike*, *The Juliet Letters*, *Kojak Variety*, *North*, *Mighty Like a Rose*, *Cruel Smile*; Shakira: *Fijación Oral Vol. 1*, *Oral Fixation Vol. 2*; Tori Amos: *Under the Pink*; Bryan Ferry: *Boys and Girls*; Kate Bush: *The Sensual World*, *This Woman's Work*; Paula Cole: *This Fire, Courage.*

KING, JACQUIRE

Kings of Leon: *Aha Shake Heartbreak, Only by the Night, Come around Sundown*; Norah Jones: *The Fall*; Tom Waits: *Blood Money, Mule Variations, Alice, Orphans: Brawlers, Bawlers and Bastards*; Modest Mouse: *Good News for People Who Like Bad News*.

KRAMER, EDDIE

Jimi Hendrix: *Are You Experienced?*, *Axis*: *Bold as Love, Electric Ladyland, Band of Gypsys*; Led Zeppelin: *Led Zeppelin II, Led Zeppelin III, How the West Was Won, Houses of the Holy, Physical Graffiti, Coda*; Traffic: *Traffic*; The Nice: *Nice, Five Bridges*; Peter Frampton: *Frampton Comes Alive!*

LANGE, ROBERT JOHN "MUTT"

AC/DC: *Highway to Hell, Back in Black, For Those about to Rock We Salute You*; Def Leppard: *High 'n' Dry, Pyromania, Hysteria, Adrenalize*; Foreigner: *4*; The Cars: *Heartbeat City*; Bryan Adams: *Waking up the Neighbours, 18 till I Die*; Shania Twain: *The Woman in Me, Come on Over, Up*; The Corrs: *In Blue*; Nickelback: *Dark Horse*; Maroon 5: *Hands All Over.*

LANOIS, DANIEL

U2: *The Joshua Tree, Achtung Baby, All That You Can't Leave Behind, No Line on the Horizon*; Peter Gabriel: *So, Us*; Bob Dylan: *Oh Mercy, Time Out of Mind*.

LECKIE, JOHN

Pink Floyd: *Meddle*; Radiohead: *The Bends*; Muse: *Showbiz, Origin of Symmetry*; The Stone Roses: *The Stone Roses*; The Verve: *A Storm in Heaven*; Kula Shaker: *K*; My Morning Jacket: *Z*; The Coral: *Butterfly House;* Cast: *All Change, Mother Nature Calls;* Doves: *Kingdom of Rust.*

LEHNING, KYLE

Randy Travis: *Storms of Life, Always & Forever, No Holdin' Back, This Is Me, Inspirational Journey, Rise and Shine, Around the Bend*; Dan Seals: *Stones, Harbinger, Rebel Heart*; Ronnie Milsap: *Heart and Soul*; Bryan White: *Bryan White, Between Now and Forever, Dowdy Ferry Road, Some Things Don't Come Easy*; Derailers: *Here Come the Derailers, Genuine*.

LETANG, RENAUD

Feist: *Let It Die, The Reminder*; Alain Souchon: *C'est déjà ça*; Peaches: *Fatherf**ker.*

LIPSON, STEVE

Annie Lennox: *Diva*; Frankie Goes to Hollywood: *Welcome of the Pleasuredome, Liverpool*; Paul McCartney: *Flowers in the Dirt*; Grace Jones: *Slave to the Rhythm*; Jordin Sparks: *Battlefield*; Boyzone: *Where We Belong, By Request*; Rachel Stevens: *Funky Dory*.

LORD-ALGE, CHRIS

James Brown: *Living in America*; Prince: *Batman* soundtrack; Joe Cocker: *Unchain My Heart*; Chaka Khan: *Destiny*; Green Day: *Nimrod, American Idiot, 21st Century Breakdown*; My Chemical Romance: *The Black Parade*; Stevie Nicks: *Trouble in Shangri-La*; POD: *Testify*; AFI: *Decemberunderground*; Darren Hayes: *Spin*; Creed: *Full Circle*; Sum 41: *Underclass Hero*; Switchfoot: *The Beautiful Letdown, Nothing Is Sound*; Slipknot: *Dead Memories, Sulfur*; Stone Temple Pilots: *Stone Temple Pilots.*

LORD-ALGE, TOM

Pink: *I'm Not Dead*, *Funhouse*; Avril Lavigne: *Let Go, The Best Damned Thing*; Blink-182: *Blink-182, Take off Your Pants and Jacket*; Steve Winwood: *Back in the High Life, Roll with It;* Live: *Throwing Copper;* Crash Test Dummies: *God Shuffled His Feet;* Hanson: *Middle of Nowhere;* The Rolling Stones: *Bridges to Babylon;* Santana: *Supernatural;* McFly: *Radio:Active;* Marilyn Manson: *The Dope Show, Mechanical Animals;* Sum 41: *All Killer No Filler, All the Good S**t;* The Fratellis: *Here We Stand;* Korn: *Life Is Peachy;* Weezer: *Weezer, Maladroit.*

MALOUF, BRIAN

Everclear: *Sparkle and Fade*; David Gray: *White Ladder*; Lit: *A Place in the Sun*; Madonna: *I'm Breathless*; Pearl Jam: *Ten*; Slaughter: *Stick It to Ya*.

MARASCIULLO, FABIAN

T-Pain: *Epiphany*; Birdman: *5 Star Stunna*; Toni Braxton: *More Than a Woman*; 50 Cent: *Before I Self Destruct*; TLC: *3D*; J.Lo: *Ain't It Funny*; Trey Songz: *Trey Day*; Monica: *After the Storm*; Lil Wayne: *Tha Carter, Tha Carter II, Tha Carter III*; Flo Rida: *Mail on Sunday*; Chris Brown: *Exclusive.*

MARROQUIN, MANNY

Kanye West: *College Dropout, Late Registration, Graduation, 808s and Heartbreaks*; Duffy: *Rockferry*; Rihanna: *Good Girl Gone Bad* (including "Umbrella"); Alicia Keys: *Songs in A Minor, Unplugged, Diary of Alicia Keys, As I Am*; Whitney Houston: *My Love Is Your Love*; Usher: *Here I Stand*; Mary Mary: *Thankful, Get Up*; John Mayer: *Continuum*; Janet Jackson: *Damita Jo*; Common: *Be*; Natasha Bedingfield: *Unwritten*; Mariah Carey: *The Emancipation of Mimi*; John Legend: *Get Lifted*; Faith Evans: *The First Lady*; Sisqo: *Return of the Dragon.*

MASERATI, TONY

Mary J. Blige: *What's the 411?*, *My Life*, *Share My World*, *Ballads*, *No More Drama*; Mariah Carey: *Daydream*, *Butterfly*; Destiny's Child: *Survivor*; Faith Evans: *Faith*, *Keep the Faith*; R. Kelly: *R*, *TP2.com*; John Legend: *Once Again*; Black Eyed Peas: *Elephunk*, *Monkey Business*; Jennifer Lopez: *J.Lo*; Alicia Keys: *Diary of Alicia Keys.*

MASSENBURG, GEORGE

Little Feat: *Feats Don't Fail Me Now, Hoy-Hoy, Let It Roll, Shake Me Up;* Linda Ronstadt: *Mas Canciones, Frenesi, Feels Like Home, Dedicated to the One I Love, We Ran;* Lyle Lovett: *Joshua Judges Ruth;* Bonnie Raitt: *Nine Lives;* Toto: *The Seventh One*; Earth, Wind, & Fire: *Gratitude, That's the Way of the World, Spirit, I Am, All N All*; Journey: *Trial by Fire, When You Love a Woman.*

MOULDER, ALAN

Elastica: *Elastica*; The Jesus and Mary Chain: *Honey's Dead*, *Automatic*; Marilyn Manson: *Portrait of an American Family*; My Bloody Valentine: *Glider*, *Loveless*, *Tremolo*; Nine Inch Nails: *The Downward Spiral*, *The Fragile*, *The Perfect Drug*; The Smashing Pumpkins: *Siamese Dream*, *Mellon Collie & the Infinite Sadness*, *Machina/The Machines of God*; U2: *Pop*.

MURPHY, SHAWN

Braveheart, Dances with Wolves, E.T. the Extra-Terrestrial, Ghost, Jurassic Park, Men in Black, Pretty Woman, Saving Private Ryan, The Sixth Sense, Star Wars: Episode I–The Phantom Menace, Episode II–Attack of the Clones, Titanic.

NICHOLS, ROGER

Steeley Dan: *Can't Buy a Thrill, Countdown to Ecstasy, Pretzel Logic, Gaucho, Aja, Two against Nature, Everything Must Go.*

NIEBANK, JUSTIN

Marty Stuart: *Country Music*; Keith Urban: *Be Here*; Patty Loveless: *On Your Way Home*, *Dreamin' My Dreams*; Vince Gill: *Next Big Thing*, *These Days.*

OLSEN, KEITH

Fleetwood Mac: *Fleetwood Mac*; Foreigner: *Double Vision*; Scorpions: *Crazy World*; Whitesnake: *Whitesnake, Slide It In, Slip of the Tongue.*

ORTON, ROBERT

Lady Gaga: *The Fame;* New Kids on the Block: *The Block;* The Police: *Certifiable;* Tatu: *200 km/h in the Wrong Lane, Dangerous and Moving;* Pet Shop

Boys: *Fundamental, Concrete;* Sean Paul: *Imperial Blaze;* Pixie Lott: *Turn It Up;* Little Boots: *Hands.*

PADGHAM, HUGH

Genesis: *Abacab*, *Genesis*, *Invisible Touch*; Phil Collins: *Face Value*, *Hello, I Must Be Going!*, *No Jacket Required*, *… But Seriously*, *Both Sides*; The Police: *Ghost in the Machine*, *Synchronicity*; Sting: *Nothing Like the Sun*, *The Soul Cages*, *Ten Summoner's Tales*, *Mercury Falling*; Peter Gabriel: *Peter Gabriel*; XTC: *Black Sea*.

PANUNZIO, THOM

U2: *Rattle & Hum*; Deep Purple: *The Battle Rages On*; Black Sabbath: *Reunion;* Ozzy Osbourne: *Live at Budokan*; Willie Nile: *Golden Down*; Jeff Healey Band: *Cover to Cover, See the Light;* Motörhead: *Hammered*; Bruce Springsteen: *Born to Run;* New Found Glory: *Coming Home.*

PARR, STEVE

Rob Lane: *John Adams* (original TV soundtrack); LTJ Bukem: *Planet Earth*; Studio Voodoo: *Club Voodoo.*

PARSONS, ALAN

The Beatles: *Abbey Road*; Pink Floyd: *Dark Side of the Moon*; Al Stewart: *Year of the Cat*; Paul McCartney: *Wild Life, Red Rose Symphony*; The Hollies: *Hollies*, "He Ain't Heavy, He's My Brother"; Ambrosia: *Somewhere I've Never Travelled*; Alan Parsons Project: *Tales of Mystery and Imagination, Pyramid, Eve, The Turn of a Friendly Card, Eye in the Sky.*

PENSADO, DAVE "HARD DRIVE"

Destiny's Child: *Survivor*; Beyoncé: *I Am Sasha Fierce,* "Check on It"; Nelly Furtado: *Loose*; Mary J. Blige: *The Breakthrough*, *Stronger with Each Tear*; Black Eyed Peas: *Behind the Front, Bridging the Gap, Monkey Business*; Christina Aguilera: *Christina Aguilera, Stripped*; Justin Timberlake: *Justified*; Mya: *Moodring*; Sugababes: *Sweet 7*; Pink: *Try This*; Pussycat Dolls: *PCD, Doll Domination*.

PLATT, TONY

Bob Marley: *Catch a Fire*, *Burnin'*; Toots & the Maytals: *Funky Kingston*; Aswad: *Aswad*; AC/DC: *Highway to Hell*, *Back in Black*; Foreigner: *4*; Boomtown Rats: *The Fine Art of Surfacing*; Anathema: *Eternity*.

POWER, STEVE

Robbie Williams: *Life thru a Lens, I've Been Expecting You, Sing When You're Winning, Swing When You're Winning, Escapology*; Babylon Zoo: *The Boy with the*

X-Ray Eyes (including "Spaceman"); Babybird: *Ugly Beautiful* (including "You're Gorgeous"); Busted: *A Present for Everyone*; Delta Goodrem: *Mistaken Identity*; McFly: *Motion in the Ocean*.

PREMIER, DJ

Jeru the Damaja: *The Sun Rises in the East, Wrath of the Math*; Group Home: *Livin' Proof*; Nas: *Illmatic, It Was Written, I Am…, Stillmatic*; The Notorious BIG: *Ready to Die, Life after Death*; Jay-Z: *Reasonable Doubt, In My Lifetime, Hard Knock Life, The Life & Times of Shawn Carter*; Rakim: *The 18th Letter, The Master*; Black Eyed Peas: *Bridging the Gap*; D'Angelo: *Voodoo*; Snoop Dogg: *Paid tha Cost to Be da Bo$$*; Christina Aguilera: *Back to Basics*; Xzibit: *Man vs Machine*; Gang Starr: *No More Mr Nice Guy, Step in the Arena, Daily Operation, Hard to Earn, Moment of Truth, The Ownerz*.

PRICE, BILL

The Sex Pistols: *Never Mind the Bollocks*; The Clash: *The Clash*, *Give 'Em Enough Rope*, *London Calling*, *Sandinista!*; The Pretenders: *Pretenders*, *Pretenders II*; Elton John: *Too Low for Zero*; Pete Townshend: *Empty Glass*; The Jesus & Mary Chain: *Darklands*; Babyshambles: *Down in Albion*.

PUIG, JACK JOSEPH

Snow Patrol: *Eyes Open*; Goo Goo Dolls: *Let Love In*; Black Eyed Peas: *Monkey Business;* Fergie: *The Dutchess;* Mary J. Blige: *The Breakthrough;* Pussy Cat Dolls: *PCD;* Stereophonics: *You've Got to Go There to Come Back;* Sheryl Crow: *C'mon C'mon;* The Rolling Stones: *Forty Licks, A Bigger Bang, Biggest Mistake;* Green Day: *Warning;* No Doubt: *Return of Saturn;* Hole: *Celebrity Skin;* Weezer: *Pinkerton;* Jellyfish: *Spilt Milk, Bellybutton*.

RAMONE, PHIL

Paul Simon: *There Goes Rhymin' Simon, Still Crazy after All These Years*; Bob Dylan: *Blood on the Tracks*; Sinead O'Connor: *Am I Not Your Girl?*; Billy Joel: *52nd Street, Glass Houses, The Nylon Curtain, The Bridge*.

RONSON, MARK

Amy Winehouse: *Back to Black*; Christina Aguilera: *Back to Basics*; Mark Ronson: *Here Comes the Fuzz*, *Version*; Lily Allen: *Alright*, *Still*; Robbie Williams: *Rudebox*.

ROSSE, ERIC

Tori Amos: *Little Earthquakes, Under the Pink*; Lisa Marie Presley: *To Whom It May Concern*; Anna Nalick: *Wreck of the Day*; Sara Bareilles: *Little Voice;* Nerina Pallot: *Fires*.

SCHEINER, ELLIOT

Steely Dan: *Aja*, *Gaucho*, *Two against Nature*; Donald Fagan: *Nightfly*; Billy Joel: *Songs in the Attic*; Fleetwood Mac: *The Dance*; Roy Orbison: *Black and White Night*; John Fogerty: *Premonition*; Van Morrison: *Moondance*; Bruce Hornsby & the Range: *The Way It Is*.

SCHILLING, ERIC

Gloria Estefan: *Into the Light*; Ricky Martin: *Sound Loaded*; Julio Iglesias: *Quelque Chose de France*; Jon Secada: *Secada*; Bacilos: *Caraluna, Sin Vergüenza*.

SCHLEICHER, CLARKE

Martina McBride: *The Time Has Come*, *Evolution*, *Emotion*, *Martina*; Big & Rich: *Comin' to Your City*, *Between Raising Hell and Amazing Grace*; Dixie Chicks: *Wide Open Spaces*, *Fly*; Pam Tillis: *Thunder and Roses*; Sara Evans: *Born to Fly*, *Restless*; Mark Chesnutt: *Savin' the Honky-Tonk*; Taylor Swift: *Taylor Swift*.

SCHMITT, AL

George Benson: *Breezin'*; Steely Dan: *Aja*, *FM (No Static at All)*; Toto: *Toto IV*; Natalie Cole: *Unforgettable*; Diana Krall: *When I Look in Your Eyes*, *The Look of Love*; Ray Charles: *Genius Loves Company*; Jefferson Airplane: *After Bathing at Baxter's*, *Crown of Creation*, *Volunteers*; Luis Miguel: *Amarte Es un Placer*.

SEAY, ED

Dolly Parton: *White Limozeen*; Martina McBride: *The Time Has Come, The Way That I Am, Wild Angels, Evolution*; Ty Herndon: *What Mattered Most*; Pam Tillis: *Put Yourself in My Place, Homeward Looking Angel*; Lionel Cartwright: *Chasin' the Sun*.

SERLETIC, MATT

Matchbox Twenty: *Yourself or Someone Like You*, *Mad Season*, *Exile on Mainstream*, *More Than You Think You Are*; Carlos Santana: *Supernatural* (including "Smooth"); Taylor Hicks: *Taylor Hicks*; Courtney Love: *America's Sweetheart*; Rob Thomas: *Cradlesong, Something to Be*; Stacie Orrico: *Stacie Orrico*; Gloriana: *Gloriana*; Collective Soul: *Hints, Allegations, and Things Left Unsaid, Collective Soul*.

SHIPLEY, MIKE

Nickelback: *Dark Horse*; Faith Hill: *Breathe*; Def Leppard: *High 'n' Dry, Pyromania, Hysteria, Adrenalize*; The Cars: *Heartbeat City*; Shania Twain: *The Woman in Me, Come on Over, Up*; The Corrs: *In Blue*; Maroon 5: *Hands All Over*.

SHOEMAKER, TRINA

Sheryl Crow: *Sheryl Crow*, *The Globe Sessions*, *C'mon C'mon*; Queens of the Stone Age: *R*; Steven Curtis Chapman: *All Things New*; Dixie Chicks: *Home.*

SIDES, ALLEN

Joni Mitchell: *Both Sides Now*; Pussy Cat Dolls: *Doll Domination*; Phil Collins: *Testify*; Alanis Morissette: "Uninvited"; Goo Goo Dolls: *Dizzy up the Girl*; *Dead Man Walking.*

SIGSWORTH, GUY

Frou Frou: *Details*; Seal: *Seal*; Bomb the Bass: *Unknown Territory*; Bjork: *Post, Homogenic, Vespertine*; Imogen Heap: *I Megaphone*; Madonna: *Music, American Life*; Britney Spears: *In the Zone, Circus*; Sugababes: *Three, Taller in More Ways*; Alanis Morissette: *Flavors of Entanglement*; Josh Groban: *Awake.*

SMITH, DON

The Rolling Stones: *Voodoo Lounge*; Ry Cooder: *Chavez Ravine*, *My Name Is Buddy*; Stevie Nicks: *Rock a Little*, *Trouble in Shangri-La*; The Tragically Hip: *Up to Here*, *Road Apples*; Tom Petty: *Long after Dark*, *Southern Accents*, *Full Moon Fever*, *The Last DJ*; Roy Orbison: *Mystery Girl*; Eurythmics: *Be Yourself Tonight.*

SMITH, FRASER T.

Craig David: *Born to Do It, Slicker Than Your Average, Trust Me*; Rachel Stevens: *Come and Get It*; Tinchy Stryder: *Catch 22, Third Strike*; Taio Cruz; *Rokstarr*; Cheryl Cole: *Three Words*; Ellie Goulding: *Lights*; Kano: *Home Sweet Home, London Town*; Beyoncé: *B'Day*; James Morrison: *Songs for You Truths for Me*; N-Dubz: *Uncle B*; Jennifer Hudson: *Jennifer Hudson*; Pixie Lott: *Turn It Up*; Chipmunk: *I Am Chipmunk.*

STAVROU, MIKE

Siouxie and the Banshees: *Join Hands, The Scream;* T.Rex: *Dandy in the Underworld.*

STENT, MARK "SPIKE"

Spice Girls: *Spice* (including "Wannabe"), *Forever*; Bjork: *Post, Homogenic, Vespertine*; Keane: *Hopes and Fears, Perfect Symmetry*; Gwen Stefani: *Love Angel Music Baby, The Sweet Escape*; Massive Attack: *Protection, Mezzanine, Heligoland*; Beyoncé: *I Am Sasha Fierce*; Lady Gaga: *The Fame Monster*; Black Eyed Peas: *Monkey Business*; Muse: *The Resistance*; Sade: *Soldier of Love*; Ciara: *Fantasy Ride*; Britney Spears: *In the Zone*; Oasis: *Standing on the Shoulder of Giants*; Janet Jackson: *Damita Jo*; Maroon 5: *It Won't Be Long before Soon*; U2: *Pop*; Goldfrapp: *Supernature, Head First.*

STONE, AL

Jamiroquai: *Return of the Space Cowboy*, *Travelling without Moving*, *Synkronized*; Daniel Bedingfield: *Gotta Get through This*; Stereo MCs: *Connected*; Bjork: *Debut*, *Post*; Turin Brakes: *The Optimist*; Lamb: *Fear of Fours*; Eagle Eye Cherry: *Sub Rosa*; Spice Girls: *Spice*.

SWANN, DARRYL

Macy Gray: *On How Life Is*, *The Id*, *The Trouble with Being Myself*, *The World Is Yours*.

SWEDIEN, BRUCE

Michael Jackson: *Off the Wall*, *Thriller*, *Bad*, *Dangerous*; Quincy Jones: *Back on the Block*, *Q's Juke Joint*; The Jacksons: *Victory*; George Benson: *Give Me the Night*; Jennifer Lopez: *This Is Me … Then*, *Rebirth*, *Brave*.

TAN, PHIL

Mariah Carey: *Daydream*, *The Emancipation of Mimi*; Rihanna: *A Girl Like Me*, *Good Girl Gone Bad*; Leona Lewis: *Spirit*; Sean Kingston: *Sean Kingston*; Ludacris: *Release Therapy*; Fergie: *The Dutchess*; Nelly: *Suit*; Gwen Stefani: *Love Angel Music Baby*, *The Sweet Escape*; Snoop Dogg: *R&G: The Masterpiece*; Usher: *My Way*, *Confessions*, *8701*; Monica: *The Boy Is Mine*; Ciara: *The Evolution*, *Fantasy Ride*; Ne-Yo: *Year of the Gentleman*; Jennifer Hudson: *Jennifer Hudson*.

THOMAS, CHRIS

Pink Floyd: *Dark Side of the Moon*; Dave Gilmour: *On an Island*; Razorlight: *Razorlight*; U2: *How to Dismantle an Atomic Bomb*; Pulp: *Different Class*, *This Is Hardcore*; INXS: *Listen Like Thieves*, *Kick*, *X*; The Pretenders: *Pretenders*, *Pretenders II*, *Learning to Crawl*; The Sex Pistols: *Never Mind the Bollocks*; Roxy Music: *For Your Pleasure*, *Stranded*, *Siren*.

TOWNSHEND, CENZO

Snow Patrol: *Eyes Open*, *A Hundred Million Suns*; Florence and the Machine: *Lungs*; Kaiser Chiefs: *Off with Their Heads*; Editors: *The Back Room*, *An End Has a Start*; U2: *No Line on the Horizon*; Bloc Party: *A Weekend in the City*; Babyshambles: *Shotter's Nation*; Late of the Pier: *Fantasy Black Channel*.

TSAI, SERGE

Shakira: *Oral Fixation Vol. 2* (including "Hips Don't Lie"), *She Wolf*; Nelly: *Suit*; Justin Bieber: *My World 2.0*; Wyclef Jean; *The Preacher's Son*; Luscious Jackson: *Electric Honey*.

VAN DER SAAG, JOCHEM

Josh Groban: *Awake*; Andrea Bocelli: *Amore*; Katherine Jenkins: *Believe*; Seal: *Soul.*

VISCONTI, TONY

T-Rex: *Electric Warrior*, *The Slider*; David Bowie: *Diamond Dogs*, *Young Americans*, *Heroes*, *Low*, *Scary Monsters*, *Heathen*, *Reality*; Iggy Pop: *The Idiot*; The Moody Blues: *The Other Side of Life*, *Sur La Mer*; Thin Lizzy: *Bad Reputation*, *Live and Dangerous*, *Black Rose.*

WALLACE, ANDY

Nirvana: *Nevermind, From the Muddy Banks of the Wishkah*; Jeff Buckley: *Grace*; Linkin Park: *Hybrid Theory, Meteora*; Slayer: *Reign in Blood, South of Heaven, Seasons in the Abyss*; Blink-182: *Blink-182*; New Model Army: *Thunder & Consolation*; Sepultura: *Arise, Chaos AD*; Bad Religion: *Stranger Than Fiction*; Sonic Youth: *Dirty*; Rage against the Machine: *Rage against the Machine, Evil Empire*; The Cult: *Electric*; Biffy Clyro: *Only Revolutions, Puzzle*; Skunk Anansie: *Post Orgasmic Chill*; Foo Fighters: *There Is Nothing Left to Lose*; System of a Down: *Toxicity, Steal This Album, Mesmerize, Hypnotize*; Slipknot: *Iowa*; Stereophonics: *Just Enough Education to Perform*; Puddle of Mudd: *Come Clean, Life on Display*; Korn: *Untouchables*; Kasabian: *Empire*; Kelly Clarkson: *My December*; Kaiser Chiefs: *Off with Their Heads*; Avenged Sevenfold: *City of Evil, Avenged Sevenfold, Nightmare;* Rancid: *And Out Come the Wolves.*

WAY, DAVE

Christina Aguilera: *Christina Aguilera* (including "Genie in a Bottle"); Spice Girls: *Spice*; TLC: *Crazysexycool*; Michelle Branch: *The Spirit Room*; Macy Gray: *On How Life Is* (including "I Try"), *The Id, The Trouble with Being Myself, Big*; Michael Jackson: *Dangerous, Blood on the Dance Floor*; Taylor Hicks: *Taylor Hicks*; India Arie: *Acoustic Soul*; Savage Garden: *Affirmation*; Shakira: *Fijación Oral Vol. 1*; Fiona Apple: *Extraordinary Machine*; Toni Braxton: *Toni Braxton*; Sheryl Crow: *Wildflower*; Pink: *Missundaztood*; Destiny's Child: *Survivor*; Jennifer Paige: *Jennifer Paige* (including "Crush"); Boyz II Men: *Coolhighharmony.*

WORLEY, PAUL

Martina McBride: *The Time Has Come*, *The Way That I Am*, *Evolution*, *Emotion*, *Martina*; Pam Tillis: *Put Yourself in My Place*, *Homeward Looking Angel*, *Thunder and Roses*; The Dixie Chicks: *Wide Open Spaces*, *Fly*; Big & Rich: *Horse of a Different Color*, *Comin' to Your City.*

WRIGHT, TOBY

Slayer: *Divine Intervention*, *Soundtrack to the Apocalypse*; Alice in Chains: *Jar of Flies*, *Alice in Chains, Unplugged*; Korn: *Follow the Leader*; Metallica: *and Justice for All*; Mötley Crüe: *Girls, Girls, Girls*; Soulfly: *Primitive*.

YOUTH

The Verve: *Urban Hymns*; Dido: *No Angel*; Embrace: *Good Will Out*; Heather Nova: *Oyster, Siren*; Crowded House: *Together Alone*; Primal Scream: *Riot City Blues*; The Fireman: *Strawberries Oceans Ships Forest, Rushes, Electric Arguments*; Killing Joke: *Revelations, Night Time, Pandemonium, Killing Joke.*

ZOOK, JOE

Modest Mouse: *We Were Dead Before the Ship Even Sank*; Sheryl Crow: *C'mon C'mon*; Courtney Love: *America's Sweetheart*; One Republic: *Dreaming Out Loud*; Rancid: *Life Won't Wait*; Brooke Fraser: *Albertine*.

APPENDIX 2
Quote References

CHAPTER 1

1 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
2 Robinson, A. (1997). Interview with Nigel Godrich. *Mix*, August.
3 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
4 Tingen, P. (2010). Inside track. *Sound on Sound*, February.
5 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
6 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
7 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
8 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
9 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
10 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

CHAPTER 2

1 Senior, M. (2009). Interview with Bruce Swedien. *Sound on Sound*, November.
2 Parsons, A. *Art & science of sound recording: Understanding & applying equalisation*. Video available from <www.artandscienceofsound.com/>
3 Katz, B. (2002). *Mastering audio: The art and the science*. Focal Press.
4 Parsons, A. *Art & science of sound recording: Understanding & applying equalisation*. Video available from <www.artandscienceofsound.com/>
5 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
6 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
7 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
8 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
9 Tingen, P. (2008). Inside track. *Sound on Sound*, November.

CHAPTER 3

1 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
2 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
3 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
4 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
5 Tingen, P. (2008). Inside track. *Sound on Sound*, September.
6 Tingen, P. (2007). Inside track. *Sound on Sound*, December.
7 Tingen, P. (2007). Inside track. *Sound on Sound*, January.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
9 Tingen, P. (2007). Inside track. *Sound on Sound*, February.

10 Tingen, P. (2007). Inside track. *Sound on Sound*, June.
11 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

CHAPTER 4

1 Elevado, R. Available from <www.gearslutz.com/board/q-russell-elevado/115987-monitoring-dragon.html>
2 Tingen, P. (2010). Inside track. *Sound on Sound*, March.
3 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
4 Elevado, R. Available from <www.gearslutz.com/board/q-russell-elevado/117088-my-mixing.html/>
5 Buskin, R. (2004). Classic tracks. *Sound on Sound*, August.
6 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
7 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
9 Elevado, R. Available from <www.gearslutz.com/board/q-russell-elevado/117226-mixing-process-time-breaks.html>
10 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
11 Tingen, P. (2008). Inside track. *Sound on Sound*, June.
12 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
13 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
14 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
15 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
16 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
17 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
18 Buskin, R. (2000). Classic tracks. *Sound on Sound*, May.
19 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
20 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
21 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
22 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
23 Tingen, P. (2008). Inside track. *Sound on Sound*, October.
24 Tingen, P. (2009). Inside track. *Sound on Sound*, July.
25 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
26 Stavrou, M. (2003). *Mixing with your mind: Closely guarded secrets of sound balance engineering revealed*. Flux Research Pty.
27 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
28 Tingen, P. (2010). Inside track. *Sound on Sound*, February.
29 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
30 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
31 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
32 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
33 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.

CHAPTER 5

1 Senior, M. (2003). Interview with Glen Ballard. *Sound on Sound*, March.
2 Tingen, P. (2010). Inside track. *Sound on Sound*, March.
3 Tingen, P. (2008). Inside track. *Sound on Sound*, November.
4 Tingen, P. (2010). Inside track. *Sound on Sound*, August.
5 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
6 Tingen, P. (2007). Inside track. *Sound on Sound*, December.
7 Daley, D. (2004). Interview with Tom Maserati. *Sound on Sound*, February.
8 Tingen, P. (2010). Inside track. *Sound on Sound*, February.
9 Tingen, P. (2009). Inside track. *Sound on Sound*, November.
10 Tingen, P. (2000). Interview with Tom Lord-Alge. *Sound on Sound*, April.
11 Tingen, P. (2008). Inside track. *Sound on Sound*, November.
12 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
13 Tingen, P. (2000). Interview with Tom Lord-Alge. *Sound on Sound*, April.
14 Tingen, P. (2009). Inside track. *Sound on Sound*, March.
15 Tingen, P. (2009). Inside track. *Sound on Sound*, April.
16 Tingen, P. (2008). Inside track. *Sound on Sound*, July.
17 Tingen, P. (1999). Interview with Spike Stent. *Sound on Sound*, January.
18 Droney, M. (1999). Interview with Mike Shipley. *Mix*, June.

CHAPTER 6

1 Daley, D. (2006). Interview with Serban Ghenea. *EQ*, May.
2 Tingen, P. (2007). Inside track. *Sound on Sound*, November.
3 Stavrou, M. (2003). *Mixing with your mind: Closely guarded secrets of sound balance engineering revealed*. Flux Research Pty.
4 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
5 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
6 Senior, M. (2000). Interview with Steve Bush. *Sound on Sound*, March.
7 Daley, D. (2003). Vocal fixes. *Sound on Sound*, October.
8 Tingen, P. (1998). Interview with Marius de Vries. *Sound on Sound*, September.

CHAPTER 7

1 Flint, T. (2001). Interview with Tony Platt. *Sound on Sound*, April.
2 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
3 Inglis, S. (2001). Interview with Pascal Gabriel. *Sound on Sound*, April.
4 Inglis, S. (2001). Interview with Pascal Gabriel. *Sound on Sound*, April.
5 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
6 Daley, D. (2006). Interview with Serban Ghenea. *EQ*, May.
7 Doyle, T. (2007). Interview with Mark Ronson. *Sound on Sound*, May.
8 Tingen, P. (2010). Inside track. *Sound on Sound*, April.
9 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
10 Daley, D. (2004). Interview with Wyclef Jean. *Sound on Sound*, July.
11 Tingen, P. (2007). Inside track. *Sound on Sound*, November.
12 Bell, M. (1998). Interview with Steve Power. *Sound on Sound*, November.
13 Buskin, R. (2007). Classic tracks. *Sound on Sound*, August.
14 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.

15 Flint, T. (2001). Interview with Marius de Vries. *Sound on Sound*, November.
16 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.

CHAPTER 8

1 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
2 Tingen, P. (2007). Inside track. *Sound on Sound*, July.
3 Tingen, P. (2010). Inside track. *Sound on Sound*, August.
4 Tingen, P. (2008). Inside track. *Sound on Sound*, March.
5 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
6 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
7 Tingen, P. (2007). Inside track. *Sound on Sound*, March.
8 Tingen, P. (2010). Inside track. *Sound on Sound*, March.
9 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
10 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
11 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
12 Tingen, P. (2000). Interview with Tom Lord-Alge. *Sound on Sound*, April.
13 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
14 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
15 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
16 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
17 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
18 Tingen, P. (2010). Inside track. *Sound on Sound*, April.
19 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
20 Nichols, R. (2006). Across the board. *Sound on Sound*, August.
21 Tingen, P. (1997). Interview with Flood. *Sound on Sound*, July.
22 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
23 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
24 Stavrou, M. (2003). *Mixing with your mind: Closely guarded secrets of sound balance engineering revealed*. Flux Research Pty.
25 Nichols, R. (2006). Across the board. *Sound on Sound*, August.
26 Izhaki, R. (2008). *Mixing audio: Concepts, practices and tools*. Focal Press.
27 Tingen, P. (2007). Inside track. *Sound on Sound*, July.
28 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
29 Buskin, R. (2007). Classic tracks. *Sound on Sound*, August.
30 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
31 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
32 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
33 Tingen, P. (2010). Inside track. *Sound on Sound*, February.
34 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.

35 Daley, D. (2006). Interview with Steve Hodge. *Sound on Sound*, November.
36 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.

CHAPTER 9

1 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
2 Owsinski, B. (2006). The mixing engineer's handbook (2nd ed.). Thomson Course Technology PTR.
3 Tingen, P. (2000). Interview with Tom Lord-Alge. *Sound on Sound*, April.
4 Tingen, P. (2005). Interview with Joe Barresi. *Sound on Sound*, July.
5 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
6 Tingen, P. (2008). Inside track. *Sound on Sound*, November.
7 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
8 Tingen, P. (2010). Inside track. *Sound on Sound*, February.

CHAPTER 10

1 Tingen, P. (2007). Inside track. *Sound on Sound*, July.
2 Bruce, B. (1999). Interview with Al Stone. *Sound on Sound*, December.
3 Tingen, P. (2007). Inside track. *Sound on Sound*, April.
4 Tingen, P. (2007). Inside track. *Sound on Sound*, November.

CHAPTER 11

1 Hatschek, K. (2005). *The golden moment: Recording secrets from the pros*. Backbeat Books.
2 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
3 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
4 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
5 Tingen, P. (2008). Inside track. *Sound on Sound*, March.
6 Tingen, P. (2008). Inside track. *Sound on Sound*, February.
7 Daley, Dan (2006). Interview with Steve Hodge. *Sound on Sound*, November.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
9 Tingen, P. (2007). Inside track. *Sound on Sound*, August.
10 Tingen, P. (2007). Inside track. *Sound on Sound*, August.
11 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
12 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
13 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
14 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
15 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
16 Tingen, P. (2009). Inside track. *Sound on Sound*, January.
17 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
18 Tingen, P. (2008). Inside track. *Sound on Sound*, April.

CHAPTER 12

1 Tingen, P. (2003). Interview with Rich Costey. *Sound on Sound*, December.
2 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
3 Daley, D. (2006). Interview with Serban Ghenea. *EQ*, May.
4 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
5 Buskin, R. (1997). Interview with Jon Gass. *Sound on Sound*, August.
6 Tingen, P. (2008). Inside track. *Sound on Sound*, December.
7 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
9 Tingen, P. (2010). Inside track. *Sound on Sound*, March.
10 Barbiero, M. (2005). Interview with Andy Wallace. *Mix*, October.
11 Tingen, P. (2007). Interview with DJ Premier. *Sound on Sound*, July.

CHAPTER 13

1 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
2 Tingen, P. (2009). Inside track. *Sound on Sound*, July.
3 Tingen, P. (2007). Inside track. *Sound on Sound*, December.
4 Buskin, R. (1997). Interview with Jon Gass. *Sound on Sound*, August.
5 Tingen, P. (2007). Inside track. *Sound on Sound*, December.
6 Tingen, P. (2009). Inside track. *Sound on Sound*, February.
7 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

CHAPTER 14

1 Buskin, R. (2004). Classic tracks. *Sound on Sound*, September.
2 Tingen, P. (2007). Inside track. *Sound on Sound*, April.
3 Tingen, P. (2007). Inside track. *Sound on Sound*, April.

CHAPTER 15

1 Tingen, P. (2007). Interview with DJ Premier. *Sound on Sound*, July.
2 Tingen, P. (2008). Inside track. *Sound on Sound*, August.
3 Doyle, T. (2009). Interview with Youth. *Sound on Sound*, March.
4 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
5 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

CHAPTER 16

1 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
2 Elevado, R. Available from <www.gearslutz.com/board/q-russell-elevado/117088-my-mixing.html/>
3 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
4 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

5 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
6 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
7 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
8 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
9 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
10 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
11 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
12 Senior, M. (2003). Interview with Glen Ballard. *Sound on Sound*, March.
13 Bruce, B. (1999). Interview with Al Stone. *Sound on Sound*, December.
14 Inglis, S. (2001). Interview with Guy Sigsworth. *Sound on Sound*, March.
15 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
16 Daley, D. (2005). Interview with Manny Marroquin. *Sound on Sound*, May.
17 Daley, D. (2005). Interview with Manny Marroquin. *Sound on Sound*, May.
18 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
19 Tingen, P. (2003). Interview with Elliot Scheiner. *Sound on Sound*, August.
20 Tingen, P. (2005). Interview with Steve Albini. *Sound on Sound*, September.

CHAPTER 17

1 Tingen, P. (1999). Interview with Spike Stent. *Sound on Sound*, January.
2 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
3 Tingen, P. (1999). Interview with Spike Stent. *Sound on Sound*, January.
4 Tingen, P. (2008). Inside track. *Sound on Sound*, January.

CHAPTER 18

1 Barbiero, M. (2005). Interview with Andy Wallace. *Mix*, October.

CHAPTER 19

1 Barbiero, M. (2005). Interview with Andy Wallace. *Mix*, October.
2 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
3 Tingen, P. (2008). Inside track. *Sound on Sound*, July.
4 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
5 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
6 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
7 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
8 Tingen, P. (1999). Interview with Spike Stent. *Sound on Sound*, January.
9 Tingen, P. (2008). Inside track. *Sound on Sound*, December.
10 Tingen, P. (2007). Inside track. *Sound on Sound*, September.
11 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
12 Anderson, J. (2007). Interview with Mike Shipley. *EQ*, November.
13 Daley, D. (2006). Interview with Steve Hodge. *Sound on Sound*, November.
14 Tingen, P. (2000). Interview with Tom Lord-Alge. *Sound on Sound*, April.
15 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
16 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
17 Tingen, P. (2006). Interview with Jim Abbiss. *Sound on Sound*, September.

18 Barbiero, M. (2005). Interview with Andy Wallace. *Mix*, October.
19 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
20 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
21 Inglis, S. (2001). Interview with Gus Dudgeon. *Sound on Sound*, July.
22 Buskin, R. (1997). Interview with Jon Gass. *Sound on Sound*, August.
23 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
24 Tingen, P. (2007). Inside track. *Sound on Sound*, May.
25 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
26 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
27 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.
28 Droney, M. (1999). Interview with Mike Shipley. *Mix*, June.
29 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
30 Lockwood, D. (1999). Interview with Bob Clearmountain. *Sound on Sound*, June.
31 Robinson, A. (1997). Interview with Nigel Godrich. *Mix*, August.
32 Owsinski, B. (2006). *The mixing engineer's handbook* (2nd ed.). Thomson Course Technology PTR.
33 Buskin, R. (2007). Interview with Gareth Jones. *Sound on Sound*, February.
34 Tingen, P. (2007). Inside track. *Sound on Sound*, February.
35 Tingen, P. (2007). Inside track. *Sound on Sound*, February.
36 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
37 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
38 Massey, H. (2009). *Behind the glass: Top record producers tell how they craft the hits* (Vol. II). Backbeat Books.
39 Tingen, P. (2010). Inside track. *Sound on Sound*, August.

CHAPTER 20

1 Flint, T. (2001). Interview with Tony Platt. *Sound on Sound*, April.
2 Droney, M. (2003). *Mix masters: Platinum engineers reveal their secrets for success*. Berklee Press.
3 Massey, H. (2000). *Behind the glass: Top record producers tell how they craft the hits* (Vol. I). Miller Freeman Books.

APPENDIX 3
Picture Credits

CHAPTER 1

Figure 1.2: Images courtesy of ADAM®, KRK Systems®, Behringer®, and M-Audio®.
Figure 1.3: Waterfall plots courtesy of Philip Newell, Keith Holland, and Julius Newell, "The Yamaha NS10M: Twenty Years a Reference Monitor. Why?" Research paper for the Institute of Acoustics. (You can find the whole article at *www.soundonsound.com/pdfs/ns10m.pdf*)
Figure 1.4: Ibid.
Figure 1.5: Ibid.

CHAPTER 2

Figure 2.3: Mobile telephone image courtesy of *www.istockphoto.com*.
Figure 2.6: Images courtesy of Avant Electronics®.
Figure 2.8: Images courtesy of Beyerdynamic® and Sony®.
Figure 4.1: Image courtesy of *www.shutterstock.com*.

This page intentionally left blank

Index

Note: Page numbers followed by "f" indicate figures.

C

D

L

M

N

Q

R

S

T

U

V

W

X

Y

Z