

CHAPTER 3

Low-End Damage Limitation

47

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.

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.

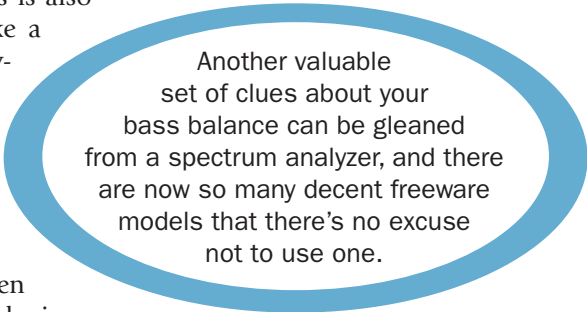


The second tactic that can help you deal with ported monitors is to block their ports, thereby defeating some of the resonance side effects. Almost any material or foam can be pressed into service here as long as it impedes the flow of air in and out of the port opening—a rolled up sock will normally 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.



Another valuable set of clues about your bass balance 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.

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 Massenburg³ 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 Chiccarelli⁴ and Eric Rosse⁵ 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.

"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

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.



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 frantically 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

A frantically flapping woofer cone (or hurricane-force gusts from a speaker port) can flag up subsonic problems even when you can't hear them.

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