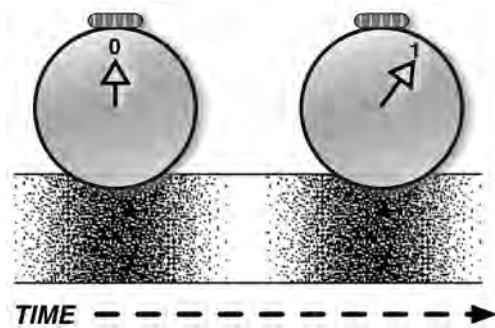


## FREQUENCY: IT MATTERS HOW OFTEN YOU DO IT

Since sound is changes of pressure, its only characteristics can be how much pressure exists at any moment, and how the pressure changes. Let’s deal with the “how it changes” part first.



**Figure 1.2** Timing from one pressure peak to the next.

Think back to the imaginary world of Figure 1.1, where we could see individual molecules. If we stand in one place, waves of pressure and rarefaction would go past us. With an imaginary stopwatch, we could measure the time from the densest part of one wave to the densest part of the next. This would tell us how quickly the tuning fork is vibrating.

Figure 1.2 shows two density peaks that are one second apart. If the vibration continues at this rate, we’d say it’s “vibrating with a frequency of one cycle per second.” That’s a mouthful, so we use the term Hertz—named after a nineteenth-century physicist—or its abbreviation Hz instead.

### Fast pressure changes are heard as sounds

In Figure 1.2, our imaginary stopwatch is vibrating at 1 Hz . . . which, incidentally, is too slow for our ears to pick up.<sup>3</sup> Another measurement—kiloHertz, or kHz—represents 1,000 cycles per second, which is friendlier to our ears.

### Frequency ranges

It’s generally accepted that humans can hear sounds in a range between 20 Hz and 20 kHz. That’s a little like saying, “Humans can run a mile in four minutes”—maybe some humans can, but I certainly can’t. A few exceptional humans can hear up to 20 kHz, but even the best hearing deteriorates when you get older. Fortunately, very few useful sounds extend to these limits. If all you consider are basic vibrations:

- The highest note of a violin is about 3.5 kHz.
- The highest note on an oboe is around 1.8 kHz.
- In fact, of all the instruments in a classical orchestra, only the pipe organ can vibrate faster than 5 kHz.

Figure 1.3 shows the basic vibration frequencies of various instruments.

<sup>3</sup> I don’t care. It was an *imaginary* sound wave.




Figure 1.3 Basic vibrations of common instruments.

from: Producing Great Sound for Film and Video, 4th ed.


© 2014 Jay Rose

## Harmonics

The fastest that a violin string or oboe reed can vibrate is considerably less than 5 kHz. But frequencies higher than that are still important. To see how, we need a closer look at the pressure waves.

As you know, a microphone converts sound vibrations into varying electrical voltages. If we connect the microphone to an oscilloscope—a device that draws a graph of voltage changes over time, such as a video waveform monitor—we can see the wave with considerable detail. Figure 1.4 shows how an ideal wave looks, photographed on my oscilloscope. **Sound 1.1** lets you hear that wave.

Positive pressure generates a positive voltage, forming the peaks. Negative pressure generates negative voltage, forming the valleys. The two voltage peaks A and B are like the two pressure peaks we measured in Figure 1.2.



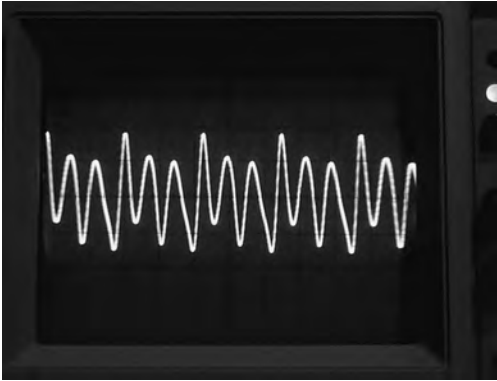
**Figure 1.4** A pure wave on an oscilloscope, with some letters and arrows added for explanation.

But this is a pure wave, also known as a *sine wave* from its mathematical function.<sup>4</sup> In the real world, most things don't have this simple, symmetrical back-and-forth vibration. Different parts of the object vibrate at different speeds, so other smaller wiggles get mixed with its main one.


The basic back-and-forth movement is called the *fundamental*. It carries most of the energy and is what we hear as the pitch of a sound. The smaller wiggles are *harmonics*. They take the form of softer, higher-frequency waves superimposed on the fundamental. The shape of an instrument, what it's made of, and how it's played determine which harmonics we hear.

Harmonics are what make the difference between two instruments playing the same pitch. You can see the difference in Figures 1.5 (an oboe) and Figure 1.6 (a violin); you'll hear it on **Sound 1.2** (oboe) and **Sound 1.3** (violin).

<sup>4</sup> This is the only time you'll see the words "mathematical function" in this book. from: *Producing Great Sound for Film and Video*, 4th ed.



**Figure 1.5** An oboe playing an A . . .




**Figure 1.6** . . . and a violin playing the same note.

Notice how every third wave of the oboe is stronger than the others. This is the fundamental frequency of its reed. The smaller peaks are a harmonic, three times as fast. The violin's fundamental is exactly as far apart as the oboe's in these photos, because they're both playing the same note. But violins have a more complicated mixture of harmonics, ranging from two to eight times higher than the fundamental. Their combination is what we hear as "a violin sound." Harmonics are the main way we tell a *toot* from a *squeak* on the same clarinet, or even an *ooh* from an *aah* by the same singer. They might be making their sounds on the same musical note, but harmonics let us tell them apart.

Just for fun, here's the same note on a trumpet: a stronger fundamental with a very complex harmonic pattern (Figure 1.7 and **Sound 1.4**).

Want proof that fundamentals make this much difference? Sounds 1.5 through 1.8 are the four waves you've just heard, but passed through a sharp filter that eliminates



**Figure 1.7** A trumpet's A shows a very different pattern.


This sample is a mere 15 pages, from a book with 549 pages... plus an additional 36-page chapter online... plus lots of downloadable audio and video exercises, tutorials, and diagnostics.

To get an idea of what else is in the book, download the Table of Contents from [www.GreatSound.info](http://www.GreatSound.info).

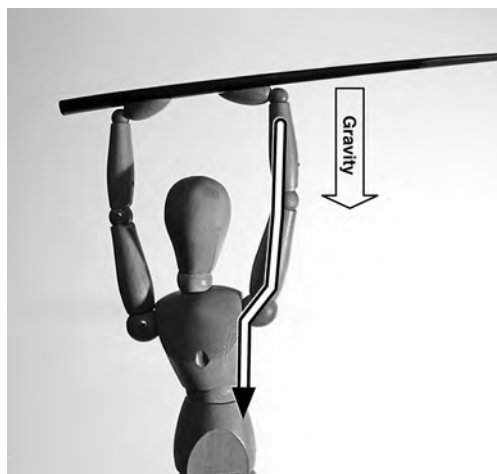
... from Chapter 7, “Microphone Technique” ...

### Boom handling

Holding a few pounds of metal and plastic over your head might seem trivial at first, but can get very heavy as the day wears on. Avoid holding the arms in a wide Y position (Figure 7.6); instead, stand like the letter H (Figure 7.7) with the arms as straight up as possible. This is a matter of physics, not aesthetics. When the arms are spread at an angle, you’re supporting the mic with your muscles . . . which will get tired. When the arms are vertical, the weight of the boom and mic is carried straight down through your bones. If you’re moving the boom around, the outside arm should be slightly bent for control (see Figure 7.9).



**Figure 7.6** Extending your arms to the sides puts the full weight of the boom on your muscles . . .



**Figure 7.7** . . . while keeping arms straight and vertical lets your bones carry most of the weight.



**Figure 7.8** Don't hold the boom like a flagpole.



**Figure 7.9** How to hold a boom properly, overhead.

It may be tempting to hold the boom like a flagpole at a parade (Figure 7.8), but this gives you less control over how the mic is aimed. Also, since the pole itself will be at an angle coming from below, it's more likely that it'll cross into a corner of the shot.

If you don't need the full length of the pole to reach the actors, hold it a foot or so in toward the middle rather than at the very end. This way part of the boom will act as a counterweight, and you won't be stressing your wrists as much.

It takes a human model to demonstrate additional details, so we'll use Christopher O'Donnell.<sup>10</sup> Note in Figure 7.9 how he keeps his inner arm—his left, in this photo—straight and rigid. It's not directly overhead, because he doesn't need the height, but vertical enough that his bones provide most of the support. The outside arm is at the end of the boom, slightly bent so it can move easily. This gives control without being tiring.

This is also a flexible position for documentary boom ops, letting you rest one end of the boom to free a hand while operating shoulder-bag equipment. Figure 7.10 shows busy production mixer G. John Garrett<sup>11</sup> boom-and-bagging a shot of a dolphin and its trainer: his right (inside) arm supports the weight straight down, while his left hand steers. He can concentrate on the shot while occasionally glancing down to read his equipment's meters.

<sup>10</sup> Not the model-turned-actor with a similar name, but an excellent feature film boom operator.

<sup>11</sup> His IMDB listing stretches back 20 years, and includes work for Disney and other major networks as well as lots of PBS documentaries.



**Figure 7.10** The same technique is adaptable while working with bagged equipment, even in slightly more challenging environments.<sup>12</sup>



**Figure 7.11** Booming at chest level: the arms still have the same functions.

wrap a couple of turns of cable around your finger where it exits the boom, so you can keep tension on the wire to stop it from swinging around inside.

You also have to secure the cable at the mic end, both to keep it taught and to provide a strain relief that protects the connector. Figure 7.12 shows the end of

If you're tall, you can hold the boom at chest level for some shots. This position is less tiring than overhead. In fact, many boom operators bring an apple box or step ladder to get some extra height, no matter how tall they are. But the strategy is still the same. Chris demonstrates in Figure 7.11: he keeps the inside arm rigid, bearing most of the weight straight down; his outside arm does the steering. Both arms are still mostly vertical.

Check the positions of the hands in those three figures as well. The inner hand is *under* the boom to act as a pivot. The outer hand is wrapped around the *top* of the boom, because the supporting hand's pivoting action makes this end push up.

In Figures 7.9 and 7.11, O'Donnell holds the cable tightly at the end of the boom instead of letting it flop. Cable management is important, because a section of wire hitting the pole will make noise. That's also why there are a couple of loops at his belt: he can move the boom suddenly, and the loops prevent cable from dropping or rubbing against the floor.

The wire spirals around the pole in the O'Donnell photos, a common practice outside of Los Angeles. Some boom designs let you run wires inside the pole, as Garrett's does in Figure 7.10; this seems to be the preferred Hollywood method. It doesn't make any difference which you choose, so long as the cable is kept tight and can't swing or bounce when you move the boom. If you're using inside wiring,

<sup>12</sup> Photo courtesy Chedd-Angier Production Company.  
from: *Producing Great Sound for Film and Video*, 4th ed.

O'Donnell's boom, with a couple of hairballs doing that job. Common practice, here in New England, is to mount the Ponytailers with "balls in back," facing the rear end of the mic. This way, you can glance at the entrance to the windscreen and see which way the mic, inside, is pointed. (You can also use tape to secure the cable.) If the cable is inside the pole, add a small loop of wire at the mic end; it protects the mic and connector if something snags the wire between takes.



**Figure 7.12** Hairballs in use. Not only do they keep things neat; they also let you know that this mic is pointing down.

### The voice, and just the voice

With the boom pointing down from above, it can be hard to avoid picking up actors' footsteps as they move around the scene.

This is a problem, since steps should be treated differently from dialog in postproduction. The techniques that make a voice crisp and understandable can make footsteps unnaturally sharp. Besides, steps are often added to or moved, depending on how a scene has been edited. Footsteps *in the clear*, while the actor is pausing, can be moved to a separate track. But those that occur during lines—if the actor is “walking and talking”—cause trouble.

- The first defense against footsteps is to get the mic as close to the actors' mouths—and as far from the floor—as possible. If you can hold the boom so its least sensitive side (usually about 120° off axis) is toward the floor, even better.
- If that's not enough to do the job, get the actors to take off their shoes. If they refuse, lay some extra carpet down for them to walk on.
- If stocking feet or extra carpet will ruin the shot,<sup>13</sup> you'll have to treat the actors' shoes. Foam-covered moleskin and similar stick-on pads from film sound suppliers can help a lot. At a minimum, get them to wear softer-soled shoes for the scene.

---

Want to read more? Full details and discount sales at [www.GreatSound.info](http://www.GreatSound.info).



... from Chapter 13 ...

## Track splitting

The main technique that separates theatrical dialog from other forms is *track splitting*. It serves two primary purposes:

- **Hiding subtle jumps in timbre and background sound:** There can be thousands of these in a film. They're inevitable when you cut from one camera angle to another, because sound changes with each camera setup: boom mics move to a different position, characters might be closer to reflective surfaces, and unless you're in a studio, the background noise keeps changing. Think about it: when you condense a morning's worth of out-of-sequence shots into a single scene, stuff is going to change . . . even if the dialog is supposed to be continuous.
- **Simplifying the mix:** The goal is a single dialog track, coming from the center speaker in the theater. But breaking it up onto a half-dozen or more faders lets you preset levels and processors for each character or mic setup, and crossfade between them as the scene progresses. This is easier and faster than trying to program all those changes into an automation system.

Figure 13.18 shows some of the track splitting in a scant 15 seconds of a fairly simple scene.

- On the bottom (letter A in Figure 13.18) is how the picture editor cut the scene in an NLE, and passed it to audio post as an OMF. This scene actually consists of a master shot and a couple of close-ups.

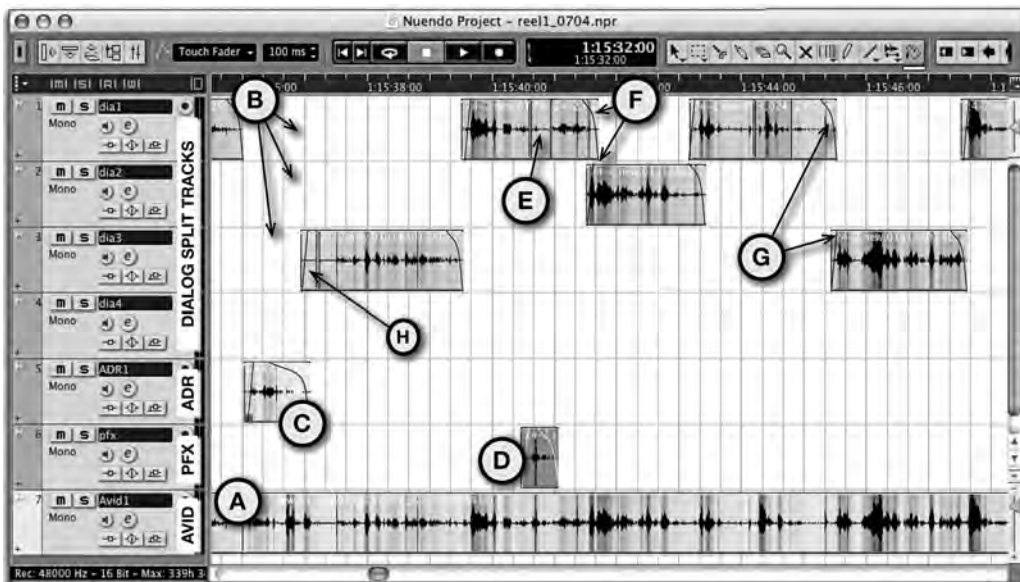


Figure 13.18 Dialog editing in 15 seconds of a simple scene.

- The scene has three different characters. Letter B shows how where individual lines have been copied from the NLE's version, and split among the top three tracks. This lets you apply different volume and equalization settings to each character, or if you want to subtly emphasize one character during the mix.
- One line had to be replaced with ADR and will need completely different processing to match, so it's separated at letter C.
- One sound effect—a door slam—was picked up by the boom during the action. It's not dialog but might be useful in the mix, so it's moved to a production effects track (D). That slam might be replaced with a different one, recorded separately. So a little bit of roomtone is copied from elsewhere in the take, and pasted where the slam was (E). This way, the background will continue if a different slam is used.
- Note how most of the clips on the top three tracks overlap by a half-second or so (such as at F), even though there are hard cuts between them on the original. These clips were extended by grabbing their handles from the OMF.
- The curved lines over the ends of most clips (G) are fades. They work in pairs as crossfades, disguising differences in roomtone. Most good audio programs will add these fades automatically.
- We were lucky mixing this scene. The actors had paused between their lines, so when the editor created overlaps, there was room ambience to cover the crossfades. But you can see (at H) how a small clip of extra roomtone had to be stuck to the front of one dialog clip. We found some clean ambience elsewhere in the take, where no one was speaking, and used it here.

## \* TIDBIT

### Why not just process the clips?

Most audio software lets you apply volume, equalization, and other processing to individual clips on a track separately, or apply them to the entire track as a whole. The former is handy for repairs and effects where just one line—or sometimes one word—needs treatment. But the latter is faster, easier to change, and usually leads to a smoother mix.


Processing individual clips locks you into their sound. Depending on the program, restoring the original can be difficult or impossible. It's seldom a good idea to make this decision until all the elements are in place, and you're monitoring in a proper environment.

If you are tempted to pre-process a clip to solve a problem, keep an unprocessed version of it on an unused track. Mute this *mirror* track so it doesn't confuse things. You'll be glad it's there, when the processing doesn't mix quite the way you wanted.

There's another feature film technique that also results in better tracks, the *dialog premix*. We'll cover that in Chapter 17.

## EDITING V: THE TRICKS

### Breathless

 Go back to **Sound 13.3**, one you loaded at the start of this chapter, or look at Figure 13.6. There's a big puff of air in the middle. Most announcers will take a gigantic breath at the start of a paragraph, and grab loud, fast ones between some sentences. These noises really say "I'm reading from a script," because nobody ever takes them during natural ad-lib conversations. They should come out of voice-over tracks.

Breaths before a paragraph are easy to get rid of: just move the in-point a little later. But those catch-breaths during a read are more troublesome. You can't just cut them out—that pulls the surrounding words too close together. You can't just erase them or turn their volume down, that leaves the words too far apart. If the announcer is on-camera, you can't change their timing at all.

Voice-over breaths can usually be replaced by pauses *two-thirds* their length. If it takes one second for the announcer to gasp, replace it with 20 frames of silence or roomtone. The result will be cleaner, more energetic, and completely natural.

I don't know why two-thirds is the magic number, but it's worked for me in thousands of projects with hundreds of different announcers. Even though I edit by ear—erasing the entire pause, playing the line in real-time, and tapping the Mark key where I think the next phrase should start—it almost always turns out to be two-thirds.

If you're starting with a clean voice-over recording and it will be used on a limited medium such as broadcast TV, or be mixed with music, you can replace breaths with digital silence. But if it'll be played at theatrical levels, use roomtone; otherwise, the finished edit will sound choppy.

On-camera breaths are trickier, since the two-thirds trick would destroy sync. Replacing the whole breath with roomtone can also be a mistake: if we see the talent's mouth open, we should hear something. Consider these on-camera breaths as you would any of the other meaningless placeholders in natural speech, like *uh* and *er*: fade them down enough to be unobtrusive, but not so much that you notice the gap. It's easier to do this as a clip edit, rather than trying to move faders quickly during the mix:

1. Split or cut the clip at the start of the breath or noise, and at the start of the following word.
2. Lower the volume of the small clip you've just isolated. A good place to start is 6 dB reduction, but it depends on the talent.
3. Make sure there are crossfades between it and the rest of the original clip.

## Simply shocking!

Amateur announcers and real-life interview subjects don't take the big breaths, but often do something worse: they unconsciously let their breath build up, then blast it out when they talk. It causes a little click if the phrase starts with a vowel. These *glottal shocks* are a natural result of nervousness in the throat, so we all get used to associating them with that condition. Your listeners won't know why, but they'll know your track is somehow hesitant and unconvincing.

Glottal shocks often follow long pauses, when a speaker is unsure and looking for "just the right word." If you try to shorten this pause but don't take care of the shock, your editing won't sound natural. It's easy to get rid of the pesky things, but only if you zoom in. Deleting about a hundredth of a second, at the start of a word, is usually enough to turn a speaker from nervous to confident.

## Extensions

When cutting documentaries or testimonials, you frequently need to start or end a voice segment on a particular syllable, even though the speaker might have more to say before or after the clip. If they didn't pause exactly where you want to edit, even the most accurate edit will sound abrupt as their voice suddenly stops or starts.

Add a little roomtone or natural background noise, butted right up to the edit. This will make it sound as though the performer—and not your editing system—made the pause.

## Snap, crackle . . .

When the clock rolls around to winter, the announcers start ticking. You can hear annoying clicks and snaps in their voices, particularly when they're close-miked and saying sounds that require the tongue to touch the roof of the mouth (like /l/ or /k/). It's caused by central heating! This dries up their mouths, thickening saliva, which starts to stick and stretch and . . . well, the sound is almost as disgusting as the description.

Unfortunately, cutting out the snap entirely can destroy the rhythm of the word. Erasing it—or replacing it with silence—leaves a noticeable hole.

But you can usually replace the snap with a tiny snippet of the vowel immediately before or after it. Locate a snap, select about the same length of audio right next to it—it'll be less than a frame—copy, and paste over the snap. Depending on the program, you may need to use a *Replace* or *Overdub* function; you want the new sound to fit over the snap, rather than move the snap later.

Because these snaps are so fast, it's almost impossible to fix them in a frame-based NLE. Use an audio program. Or hand the announcer a glass of water, and re-record.

## Squeezing and stretching

Modern DAWs let you change the timing of a clip, speeding or slowing words without affecting their pitch or timbre. Obviously, it can only be done during cutaways, when

Full details and discount sales at [www.GreatSound.info](http://www.GreatSound.info).

---

... from Chapter 16, “Processing” ...

## Getting the most from an equalizer

### Learn the bands

The best tool for learning how to equalize is the equalizer itself. Grab a short voice clip, such as **Sound 13.2**, and set it to play continuously. Start with all of the level controls in their neutral position (usually marked 0 dB) and the Q around 7. Pick a section, raise its level about 6 dB, and listen to the audio while you sweep the frequency control very slowly. The changes will be subtle as you move from one frequency to the next, but try to give a name to what you hear: *boomy*, *powerful*, *harsh*, *bright*—whatever comes to mind. Then turn the section’s level as far down as it’ll go, and sweep through its frequencies again. Lowering a few frequencies can make some signals sound better. Then try the whole thing again with some music (**Sounds 14.3** through **14.6**). You’ll hear different effects depending on the type of material.

Ears get used to equalization very quickly, so keep going back to the unequalized sound as a reference, and don’t spend more than 10 or 15 minutes at a time on this exercise. Do this a few times, and you’ll start to learn exactly where to set those equalizer knobs for the sound you want.

### TRY THIS

#### Another way to learn

I’ve prepared a Web page that teaches you to identify frequency ranges. Sample male and female voices, and music in various genres, are broken into octave-wide bands. Each band is explored with spectrograms—graphic representations of time and frequency—and specially prepared sound files.

Go to [www.dplay.com/tutorial/bands](http://www.dplay.com/tutorial/bands).

## Overequalization

Once you learn how to hear subtle differences, you can avoid the most common equalizer mistake: setting the knobs too high. Equalizers are volume controls, and too much volume causes problems—particularly in digital systems. Check volumes along the entire equalized track. A 12 dB boost at 5 kHz may help the strings at the start of a piece of music, but it'll cause an awful clatter when the cymbal comes in. (Far better to use a 6 dB boost around 10 kHz.)

- As a general rule, don't raise any control higher than 6 dB—that's a one-bit boost, in digital terms—and remember that overlapping bands have their levels added together. Save the more radical settings for special effects, or to rescue badly recorded tracks.
- You usually shouldn't turn all the knobs in the same direction. If everything's boosted, the signal won't sound better—just more distorted.

Unless you're fixing specific noises at the extremes of the band, almost all equalization takes place between 200 Hz and 10 kHz . . . with most of it under 3 kHz. If you're equalizing much beyond these frequencies, you're probably adding noise.

## Equalizer tips

If you're in a hurry, these ideas can help you get started. But don't treat the settings as gospel. Every track is slightly different, and every equalizer contributes its own sound.

- **Strengthen an announcer.** Cut off everything below 90 Hz—those frequencies are just wasting power. Then try a gentle peak (3 dB,  $Q = 7$ ) around 240 Hz for warmth, and a similar boost around 1.8 kHz for intelligibility. A sharp dip (–18 dB,  $Q = 100$ ) around 5 kHz can help sibilance.
- **Help fix muddy dialog** with a cutoff below 150 Hz, and a 3–6 dB boost ( $Q = 7$ ) around 2 kHz.

### ✓ TIP

#### Turn the knobs counterclockwise!

Remember, equalizers can be turned *down* as well as up. If a voice is getting lost under the music track, don't look for a way to boost the voice. Instead, dip the music a few decibels, around 1.5–2 kHz.

Doing it this way lets the music show through without changing the voice timbre, so the overall mix continues to sound natural.

- **Make music more intense.** If there's no singing or spoken words that'll be affected by the equalizer, try boosting the bass notes (6 dB, Q = 3) around 100 Hz, and add a 6 dB high-frequency shelf around 3 kHz.
- **Simulate a telephone conversation** with the setup in Figure 16.7.
- **Tune out hum or whistles** with multiple sections of a parametric. Turn on just the lowest frequency section, and tune it using the technique explained under "Peaking equalizers." Then turn on the next higher section and tune it the same way, starting around twice the frequency you found for the first. Continue until the sound is fixed or you run out of sections.
- **Reduce distortion.** If a natural sound doesn't have many harmonics (Chapter 1), you can often clean up bad recordings by applying a sharp high-frequency cutoff filter. This not only eliminates hiss; it also reduces the artificial harmonics that a bad analog recording can generate. Start with the equalizer as low as 5 kHz, and then slowly raise its cutoff frequency until the recording doesn't seem muddy.

## COMPRESSORS

In the real world, the difference between loud and soft adds excitement to what we hear. But in the electronic world of a film or video track, loud causes distortion and soft gets lost in electronic noise. Used properly, a compressor can control those pesky level changes—and make a track sound louder—without affecting its dynamic feel. But used the wrong way, a compressor turns tracks into bland, unlistenable mush.

It's easy to misuse these things. That's because the best compressors have a daunting array of knobs with relatively non-intuitive names. Experts need this much control to properly shape the sound, because every element has its own dynamic footprint and should be handled differently. When you finish this chapter, you'll be on the way to expertise . . . and will know how to reach for the right knob every time.

### What happens inside a compressor

You can understand how a compressor works with a couple of simple graphs. There's no calculating or complicated numbers involved, just tracing your finger up a line on the page or screen.

#### Compressor input and output

The graphs starting with Figure 16.9 show how volumes can be affected by a compressor. Sounds coming in are displayed on the vertical axis: loud is at the top, soft is at the bottom. The signal travels to the black diagonal line, and reflects down to the horizontal axis. That horizontal line shows the output: soft on the left, loud on the right.

**Want to read more? Full details and discount sales at [www.GreatSound.info](http://www.GreatSound.info).**