The Short Delay, Part I:
Untangling the Comb Filter

BY ALEX CASE

The humble delay is a powerful production tool. You see, not all delays sound alike. Long delays sound very different from short delays.

"No duh, Case," you think to yourself. Let me explain. The sonic difference between a long delay and a short delay isn't just the apparent length of the delay.

Long delays are pretty intuitive; they sound like an echo, perhaps repeating a few times. Short delays, on the other hand, aren't heard as echoes. Very short delays have an important spectral effect on the sound. Then there are the delay times in between long and short. They have a more complex, textured effect.

So we classify delays into three broad categories, cleverly called long (greater than about 50 ms), medium (between about 50 ms and 20 ms), and short (less than about 20 ms). We covered long delays in last month's column; medium and short delays are so darn cool that we'll dedicate this and next months' columns to them.

Make it short

As delay times fall below about 50 milliseconds, they take on a new persona. If you are actually reading this in your studio, try the following experiment. (Those of you reading this on an airplane or tour bus are out of luck. That'll teach you: never leave your studio, ever.)

Patch up a sampler loaded with a variety of sounds or find a multitrack tape with a good variety of tracks. On your mixer/DAW, set up a delay fed by an aux send that returns to your monitor mix at about the same volume as the synth or original tracks. Pan both the source audio and the return from the delay dead center.

Listen carefully to the mix of each source sound when combined with the output of the delay. Start with a bass line. Check out the combination of the bass with a long delay, maybe 200 milliseconds. Yuk. It's a blurry, chaotic mess.

Now start shortening the delay. 100 ms, 80 ms, 60ms, 20 ms, 10 ms, 5 ms, down to 3 ms and below. Listen carefully as you do this. What the heck is going on?

The long delay is just an echo. The very short delays (15 ms and lower) sound strange, sometimes hollow, sometimes boomy. At one short delay setting there's extra low end, then at a slightly different delay time, a lack of low end. This mix of a bass sound with a very short delay sounds like it's been equalized.

Gradually lengthen the delay time and listen for the point at which it starts to sound like a distinct echo again. Depending on the bass sound, you may hear the delay separate from the bass into an echo somewhere between about 60 and 80 milliseconds. In between the very long and the very short delay times, well, it's hard to describe.

Next try a snare sound. Again start with a long delay and gradually pull it down to a short delay. Again we find it is a distinct echo at long settings. The delay introduces a strange timbral change at short delays and something tough to describe as it transitions between the two. While we're here, do the same experiment with an acoustic or electric guitar track, or a string patch on the sampler.

Welcome to the real world of delays. They aren't just for echoes anymore. When delays become shorter than about 50 or 60 milliseconds (depending on the type of sound you are listening to, as demonstrated above) they are no longer repeats or echoes of the sound. The same device that delays a signal starts to change the color, the spectral content of the signal.

Let's check out how it works.

Sine of the times

Consider first a pure tone (no fun to listen to, but helpful to study). Mixing together—at the same volume and pan position—the original signal with a delayed version of itself might have results like the two special cases shown in Figure 1. (Just look at the solid lines for now; we'll come back to the dashed lines and what they mean in a minute.)

If the delay time happens to be exactly the same as the period of the sine wave, we have the constructive interference shown in Figure 1a. That is, if the delay time we set up on our delay processor is exactly equal to the time it takes the sinusoid to go through one cycle, then they combine cooperatively, and the result is a signal of the same frequency but with twice the amplitude.

The situation in Figure 1b represents another special case. If the delay time happens to be equal to half a period (half the time it takes the sine wave to complete exactly one cycle), then the original sound and the delayed sound move in opposition to each other—they are 180 degrees out of phase. The combination results in zero amplitude—pure silence.

If you have access to a sine wave oscillator (either as test equipment or within your synthesizers or computer), give it a try. I recommend 500 Hz as a starting point—it isn't quite as piercing as the standard test tone of 1000 Hz, and the math is easy. The time it takes a pure 500 Hz tone to complete one cycle is 2 milliseconds (Period = 1 / Frequency = 1 / 500 = 0.002 seconds = 2 milliseconds).

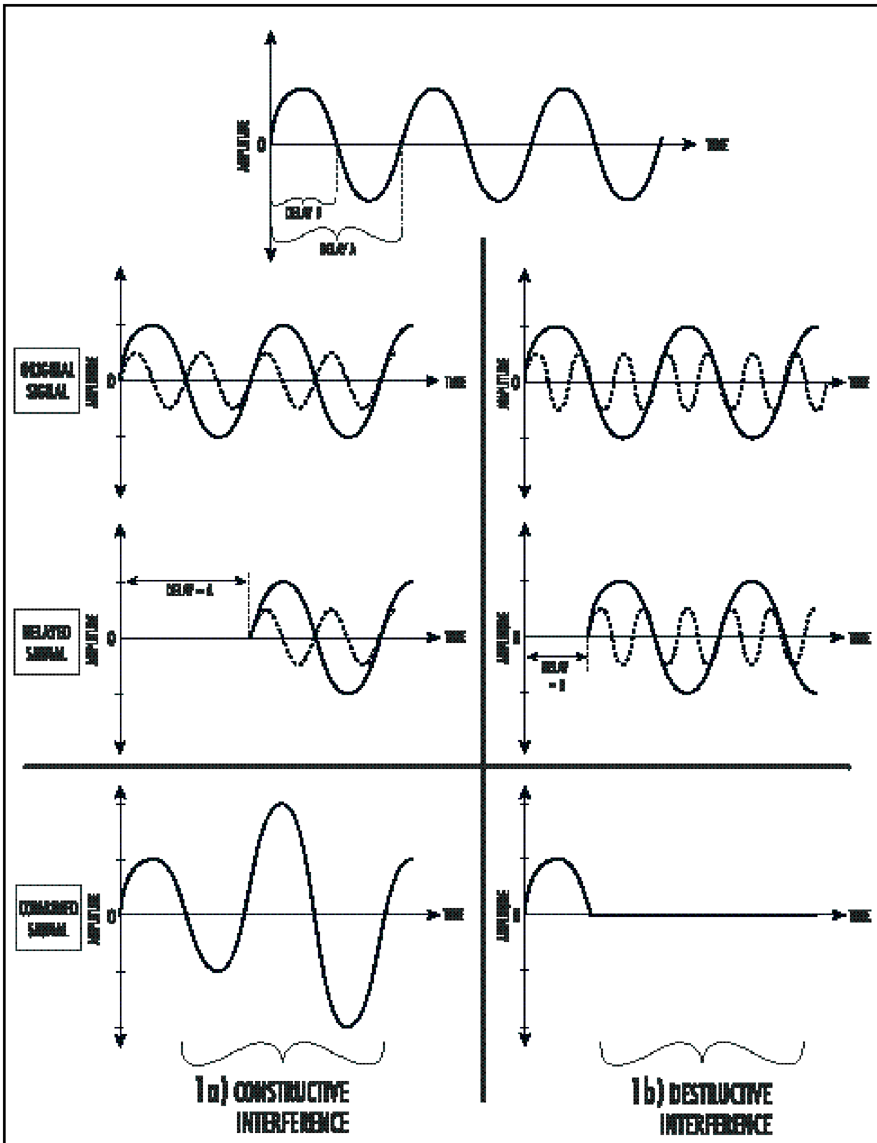

Figure 1: Combining Sine Waves

So mixing together equal amounts of the original sine wave and a 2 millisecond delayed version will create the case shown in Figure 1a. Set the delay to 1 millisecond, creating the situation of Figure 1b, and you'll find that the sine wave is essentially cancelled.

Now look at the dashed-line wave forms on Figure 1. They show that these doublings and cancellations happen at certain other higher frequencies as well. For any given delay time, certain frequencies line up just right for perfect constructive or destructive interference.

The math works out as follows. For a given delay time (t expressed in seconds, not milliseconds) the frequencies that double are described by an infinite series: 1/t, 2/t, 3/t, ... The frequencies that cancel are: 1/2t, 3/2t, 5/2t, ...

Using these equations we confirm that a 1 millisecond delay (t = 0.001 seconds) has peaks at 1000 Hz, 2000 Hz, 3000 Hz, ... and nulls at 500 Hz, 1500 Hz, 2500 Hz, ... This is consistent with our observations in Figure 1b of how a 1 millisecond delay cancels a 500 Hz sine wave.

In Figure 1a, the dashed line is the 2/t (constructive) case, and in 1b, the dashed line is the 3/2t (destructive) case.Again, you can see how the peaks and dips in the waves either add up or cancel out.

A 2 millisecond delay has amplitude peaks at 500 Hz, 1000 Hz, 1500 Hz, ... and nulls at 250 Hz, 750 Hz, 1250 Hz, ... We looked at the results of this 2 ms delay for the single frequency of 500 Hz in Figure 1a. The math reveals that the peaks and dips happen at several frequencies, not just one. Of course, the only relevant peaks and valleys are those that fall within the audible spectrum from about 20 Hz to 20,000 Hz.

To explore this further, return to your mixer setup combining a sine wave with a delayed version of itself set to the same amplitude. Sweep the sine wave frequency higher and lower, watch your meters, and listen carefully. With the delay fixed to 1 millisecond, for example, sweep the frequency of the sine wave up slowly beginning with about 250 Hz.

You'll hear the combination of the delayed and undelayed waves disappear at 500 Hz, reach a peak at 1000 Hz, disappear again at 1500 Hz, reach a peak again at 2000 Hz, and so on. We've got a delay (not an equalizer) changing the frequency content of our signals. We've got a delay (not a fader or a compressor) changing the loudness of our mix.

Let's ride the faders in the following experiment. On your mixer, one fader has the original sine wave at 500 Hz panned to center. And the sine wave is also sent to a delay unit set to a delay time of 1 millisecond. Another fader controls the return from this delay, also panned to center.

Start with both faders down. Raise the fader of the source signal to a reasonable level. Now raise the second fader. As you make the delayed signal louder, your mix of the two waves gets quieter. As you add more of the delayed sine wave, you get more attenuation of the original sine wave.

This is the phenomenon shown in Figure 1b. And the mix reaches its minimum level when the two signals are at equal amplitude.

Time for music

Stupid parlor trick or valuable music production tool? To answer this question we have to get rid of the pure tone (which pretty much never happens in pop music) and hook up an electric guitar (which pretty much always happens in pop music).

Run a guitar signal—live, from your sampler, or off tape—through the same setup above. With the delayed and undelayed signals set to the same amplitude, listen to what happens.

Can you find a delay time setting that will enable you to completely cancel the guitar sound? Nope. The guitar isn't a pure tone (thank God). It is a complex signal, rich with sound energy at a range of frequencies. No single delay time can cancel out all the frequencies at once.

But mixing together the guitar sound with a 1 millisecond delayed version of the guitar sound definitely does do something, and what happens is definitely cool. It would be nice to understand what is going on.

We already saw a 1 millisecond delay remove the 500 Hz sine wave entirely. In fact, it will do the same thing with guitar (or piano, or didgeridoo, or anything). Musical instruments containing a 500 Hz component within their overall sound will be affected by the short 1 millisecond delay; the 500 Hz portion of their sound can in fact be cancelled. What remains is the tone of the instrument without any sound at 500 Hz.

But wait, there's more. Try the 2 millisecond delay. In the case of the 500 Hz sine wave, we saw that the signal got louder when we added this delay. In the case of the guitar, the 500 Hz portion of the signal gets louder.

Taking a complex sound like guitar, which has sound energy at a vast range of different frequencies, and mixing in a delayed version of itself at the same amplitude will cut certain frequencies and boost others. This is called comb filtering (see Figure 2) because the alteration in the frequency content of the signal looks like teeth on a comb.

Combining a musical waveform with a short delayed version of itself radically modifies the frequency content of the signal. Some frequencies are cancelled, others are doubled. The intermediate frequencies experience something in between outright cancellation and full-on doubling.

So short delays are less like echoes and more like equalizers; they are too short to be perceived as echoes. In fact they are so short that they start to interact with discreet compo-

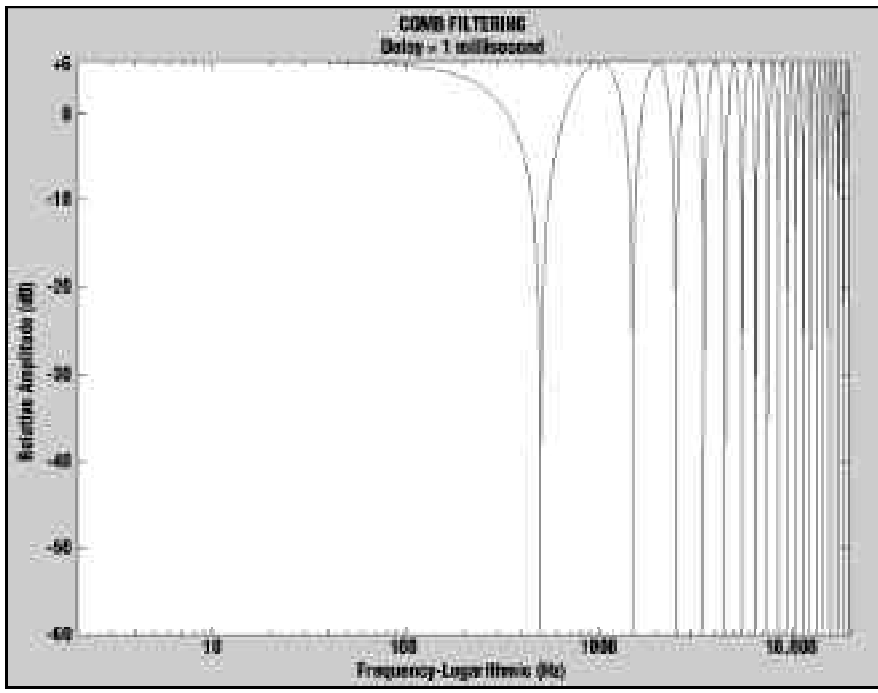

Figure 2a

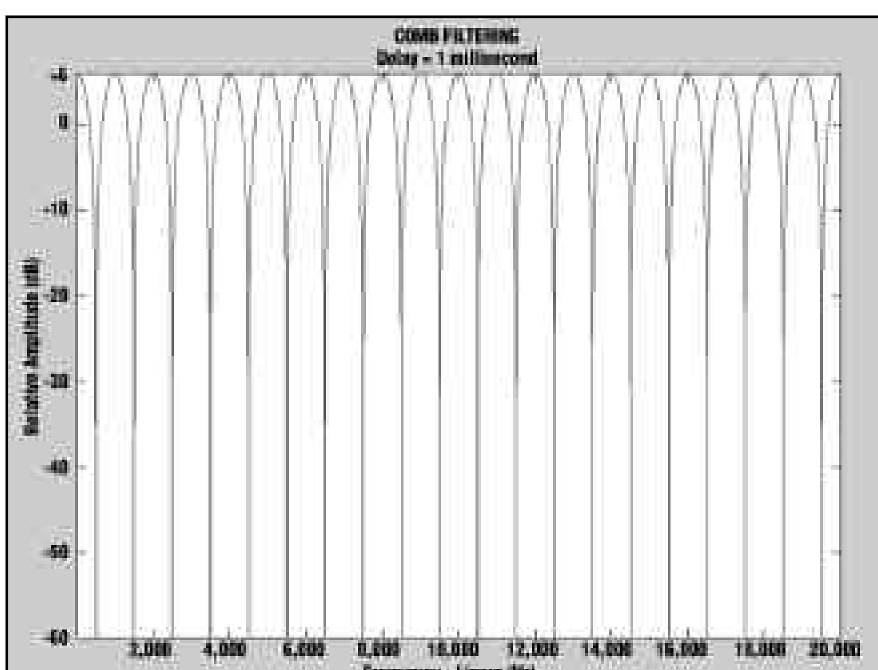

Figure 2b

NUTS and BOLTS

nents of the overall sound, adding some degree of constructive (i.e. additive) or destructive (i.e. subtractive) interference to different frequencies within the overall sound.

Figure 1 demonstrates this for a sine wave. Figure 2 summarizes what happens in the case of a complex wave like guitar, piano, saxophone, vocal, etc.

The spectral result is that the combining of a signal with a delayed version of itself acts like a radical equalization move: a boost here, a cut there, another boost here, another cut there, and so on. In theory you could simulate comb filtering with an equalizer, dialing in carefully the appropriate boosts and cuts.

That's the theory. In fact, one rarely could. To fully imitate the comb filter effect that a 1 millisecond delay creates, you'd need an equalizer with about 40 bands of eq (20 cuts and 20 boosts within the audible spectrum). I've never seen so crazy an equalizer (other than in software).

In fact, part of the point of using short delays in your mix is to create sounds that you can't create with an equalizer. It's pretty impressive. A single short delay creates a wildly complex eq contour.

Short delays offer a very interesting extra detail: they create mathematical—not necessarily musical—changes to the sound.

Study Figure 2, comparing part 2a to part 2b. They show the same information. But Figure 2a presents the information with a logarithmic frequency axis. This is the typical way of viewing music, because it's how our ears hear: double the frequency, go up an octave. Double it again, go up another octave, and so on. This relationship is why, for example, you go up a half step with each fret on a guitar but the frets get closer together as you go up the neck.

But if you look at comb filtering with a linear (and non-musical) frequency axis, you see that the peaks and dips in the filter are spaced perfectly evenly. It isn't until you view the implications of the short delay in this linear way (Figure 2b) that you see why it is in fact called a comb filter. You'll get a better hairdo using the comb in Figure 2b instead of 2a.

This highlights another unique feature of using short delays to shape the harmonic content of a sound. The distribution of the cuts and boosts is a mathematical peculiarity, all equally spaced in terms of the linear number of Hz. It is distinctly non-musical.

Patch up the comb filter with a special radical effect in mind. If you want more careful tailoring of sound, use an equalizer with its logarithmic, more musical controls.

Time for reflection

It's still fair to ask: why all this talk about short delays and their effect on a signal? After all, how often do we use delays in this way?

It is essential to understand the sonic implications of these short delays because all too often they simply can't be avoided.

Fortunately the sound reflected off the floor will also be a little quieter, reducing the comb filter effect. If the floor is carpeted, the comb filtering is a little less pronounced. Place absorption at the point of the reflection, and the comb filtering is even less audible.

An important part of the recording craft is learning to minimize the audible magnitude of these reflections by taking advantage of room acoustics in placing musical instruments in the studio and strategically placing absorptive materials around the musical source. This is one approach to capturing a nice sound at the microphone.

Better yet, learn to use these reflections and the comb filtering they introduce on purpose. For example, raising the microphone will make the difference in length between the reflected path and the direct path even longer. Raising the microphone therefore lengthens the acoustic delay time difference between the direct sound and the reflected sound, thereby changing the spectral locations of the peaks and valleys of the comb filter effect.

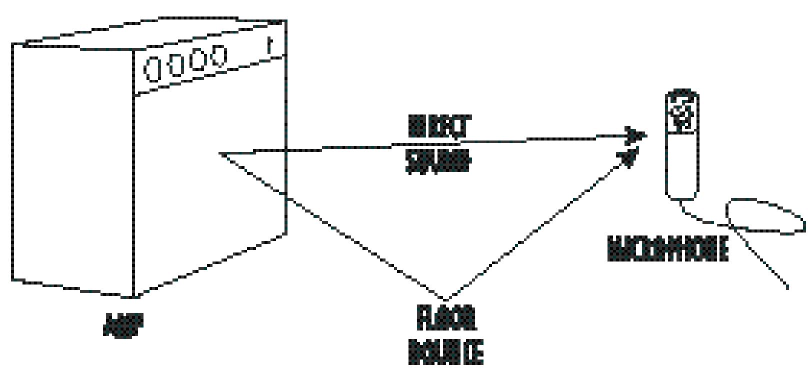

Figure 3a: Reflections Cause Comb Filtering

Consider recording an electric guitar. With the amp in the middle of the room on a beautiful wooden floor, we place a sweet tube microphone a few feet away and try to capture the natural sound of the amp in the room. This is a good approach, shown in Figure 3a.

Problem is, the sound reflected off the floor and into the microphone will arrive a split second later than the sound that went straight from amp to mike. The path is longer via the reflected path, introducing some delay.

The result is some amount of comb filtering. Recording a sound and a single reflection of that sound is a lot like mixing a track or sample with a delayed version of itself, as in our discussions above. Comb filtering is a part of everyday recording.

Of course, raising the microphone also pushes the microphone further off-axis to the amp, changing the timbre of the electric guitar tone as picked up by the microphone. You can raise the amp up off the floor, perhaps setting it on a piano bench. You can tilt the amp back so that it faces up toward the raised microphone.

But then again, you can flop the amp on its belly, facing straight down into the floor if that sounds good.Always do what sounds good. Delay-induced comb filtering is only part of the equation.

Another common approach to recording a guitar amp (and pretty much any other instrument) is to use a combination of two or more microphones to create the sound as you record it onto a single track.

Consider the session shown in Figure 3b: two microphones, one track. Here we have a close microphone (probably a dynamic) getting the in yer face gritty tone of the amp and a distant microphone capturing some of the liveness and ambience of the room. You might label the fader controlling the close microphone something like "close" and the fader governing the more distant mic something like "room." You adjust the two faders to get the right mix of close and room sounds and print that to a single track of the multitrack.

That's only half the story. As you adjust the faders controlling these two microphones, you not only change the close/ambient mix, you also control the amount of comb filtering introduced into the guitar tone. These two mics pick up very similar signals, but at different times. In other words, they act very much like the signal plus delay scenario we've been discussing.

Moving the distant microphone to a slightly different location is just like changing the time setting on the del

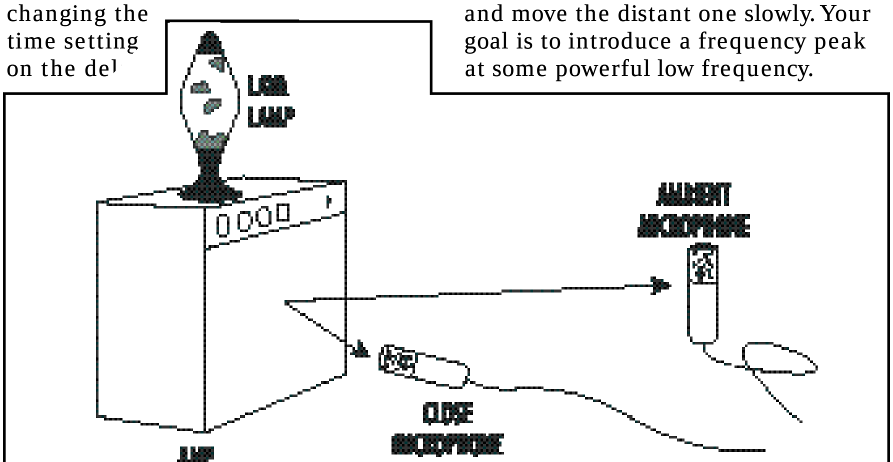

Figure 3b: Multiple Microphones Lead To Comb Filtering

unit. It effectively selects different key frequencies for cutting and boosting using the exact same principles we explored in Figures 1 and 2.

Sound travels a little farther than a foot per millisecond. To lengthen the delay time difference by about a millisecond, move the distant mic back about a foot. To get a ten millisecond delay increase, move the distant mic back about ten feet. It's that simple.

Naturally, there's too much to keep track of. Each of these microphones receives reflected sounds from the floor, the ceiling, and all the other room boundaries—all in addition to the obvious direct sound from the amp.

So we get a complex interaction of the many components of guitar sound radiating out of the amp. The direct sounds into multiple microphones arrive at different times, leading to some amount of comb filtering. The reflections from the various room boundaries into each microphone arrive at a later time than the direct sound, adding still more comb filtering.

There is an infinite number of variables in recording. In theory, we recording engineers like this complexity. (For certainty, become a tax accountant. For endless opportunities of exploration, become a recording musician.) Understanding comb filtering is part of how we master the vast recording process.

The myth of the sweet spot

Perhaps you want a tough, heavy, larger than large guitar tone. Maybe a comb filter derived hump at 80 Hz is the ticket. Or should it be 60 Hz?

You decide. Explore this issue by moving the microphones around. Place two microphones on the amp as shown in Figure 3b. Keep the close mic fixed and move the distant one slowly. Your goal is to introduce a frequency peak at some powerful low frequency.

If you have the luxury of an assistant engineer, have him or her slowly move the microphone around while you listen to the combined close/distant microphone mix. If you lack an assistant, record a take onto tape while you slowly move the microphone, as quietly as you can.

When those comb filter peaks and notches fall into frequency ranges that complement the tone screaming out of the guitar amp, you'll have found a sweet spot. No dumb luck. No magic. Finding the mic placement that captures the tone that pleases the guitarist simply requires a bit of patience—and an understanding of the spectral implications of short delays.

The art of microphone placement requires mastery of room acoustics, musical acoustics, and psychoacoustics. To achieve predictably good sounding results you need recording experience, an understanding of microphone technologies, knowledge of microphone sound qualities, exposure to the various stereo miking techniques, and many other topics.

In other words, you need a subscription to *Recording*. And an essential tool in mic placement is the use of comb filtering for fun and profit. Avoid it as necessary. Or use it on purpose when you can.

Electric guitar, which my mom and some scientists would classify as broadband noise, responds well to comb filtering. With energy across a range of frequencies, the peaks and dips of comb filtering offer a distinct, audible sound property to be manipulated.

Other instruments reward this kind of experimenting. Try placing a second (or third, or fourth…) microphone on acoustic guitar, piano, anything. Experiment with the comb filter-derived signal processing to get a sound that is natural—or wacky.

One day you may find yourself in a predicament: the amp sounds phat out in the live room, but thin in the control room. Perhaps the problem is that, courtesy of the short delay between two microphones, you've got a big dip in frequency right at a key low frequency region. Undo the problem by changing the spectral location of the frequency notch: move a microphone, which changes the delay, which changes the frequencies being cancelled.

Every time you record with more than a single microphone, make it part of your routine to listen for the comb filter effect. Check out each mic alone. Then combine them, looking for critical changes in the timbre.

What frequency ranges disappear? What frequency ranges get louder? The hope is to find a way to get rid of unwanted or less interesting parts of the sound while emphasizing the more unique and more appealing components of the sound.

And make short delays part of your mixing bag of tricks. For subtle tone shaping or a radical special effect, the short delay is a powerful signal processor. Mastering it will lead directly to better sounding recordings.

Alex Case knows the diffence between a comb filter and an oil filter.
Request Nuts & Bolts topics via case@recordingmag.com.