Bob Moog

sound sampling instruments, part 3: quick listening tests for digital samplers

LAST MONTH, WE delved into the technical side of sound sampling. We saw that sampling in a digital instrument consists of taking a rapid series of snapshots (samples) of a sound waveform, storing the snapshots as digital numbers, and playing them back to reconstruct the original waveform. In order to sample a sound, the samples must be taken at at least twice the rate of the sound's bandwidth. Noise and distortion in a sampled sound are related to the resolution (number of bits per sample). If each sample is eight bits, then the sample's noise and distortion can be no better than 48dB below the desired signal; for every extra bit of resolution, the theoretical best signal-to-noise ratio gets better by 6dB.

It's important to understand that the above signal-to-noise figures are theoretical 'best-case' numbers. Real-world sound sampling instruments generally produce more noise and distortion than the best-case numbers would lead you to believe. The reasons for this extra noise have to do with details of an instrument's hardware. For instance, sampling instruments that are built with only one sound channel generally produce more noise than instruments that are built with a separate channel for each voice. This means that two sampling instruments that both use eight-bit sampling may sound distinctly different from each other, or that an instrument with ten-bit sampling may seem to have higher fidelity than another with twelve-or even sixteen-bit sampling.

As a musician who is interested in comparing the sampling sound quality of several instruments, you should know a couple of simple listening tests that you can perform in a few minutes, without a lot of equipment, which rely primarily on your ears. Sampling, and then listening to, a conventional musical sound (such as a synthesizer tone or a cymbal crash) will tell you something about the quality of the sampler. Noise and distortion in the sample will usually cause the output to be dull and muddy. However, it is possible to construct sounds that more clearly show how much and what kind of noise and distortion are being produced.

In order to run these tests, you will need some sort of small synthesizer that is capable of producing sine-like as well as buzzy pitched tones, and of imparting slow frequency modulation (a siren effect from a sine or triangle wave LFO) and amplitude contouring (enveloping) to the tone. Of course, the instrument that you use to generate the test tones should not produce audible distortion or noise of its own. A small modular or monophonic analog synthesizer is ideal.

Test No. 1. Set up the test synth to produce a single sustained tone, of a frequency around 1kHz (three octaves above Middle C), and of low harmonic content. If your instrument produces a sine wave, use that; if not, then use any waveform and set the lowpass filter so that the overtones are filtered out, letting only the fundamental frequency through. Your test tone should be pure and flutelike, with no audible noise or distortion.

Once you've set up your test tone, sample a couple of seconds of the tone with the sampling instrument you are testing, following the manufacturer's instructions to the letter. Pay particular attention to the input level setting, to insure that your test tone is loud enough to give a good sample but not so loud as to overload the input electronics of the sampler.

Now connect the output of the instrument under test to one input of your monitor system, and the output of the test-tone synth to another input. By switching back and forth, compare the sampled sound to the original. You will clearly hear the difference between the two. Since the test tone has only one frequency component, the difference that you hear cannot be attributed to the frequency response of the instrument under test. The difference has to come from added noise and distortion introduced by the sampling process.

Test No. 2. To hear what the noise and distortion components are in some detail, repeat Test No. 1 with one additional element. This time, frequency-modulate the test tone with a slow (one cycle every second or two) sine, triangle, or sawtooth LFO. The pitch of the test tone should go up and down about an octave.

When you play back the sampled sound and A/B it with the original, you'll hear three types of additional stuff that aren't in the test tone: additional harmonics that just brighten the tone, 'whistles' that don't follow the pitch of the test tone, and hissing or rumbling that seems not to change as the test tone's pitch rises and falls. Let's look at these one at a time.

Increased brightness in a tone that is otherwise clean is a sign of *harmonic distortion*. This is the least bothersome form of dirt, since the distortion is harmonically related to the desired pitch. For nearly all traditional musical tones, a small amount of harmonic distortion is usually not noticeable. However, occasionally you'll hear a high-pitched buzz or whistle that perfectly tracks the test tone's pitch. This is called high-order harmonic distortion, and it is objectionable, especially when flutelike and similar tones of low harmonic content are sampled.

Whistles that do not track the test tone's pitch are *alias components*—interactions between the test tone and the sampling frequency that are not adequately filtered out. Alias tones are perhaps the most objectionable type of garbage that a sampling instrument can produce. Alias tones are heard as distinct pitches when pitched sounds are sampled, or as an overall muddiness when percussive or noisy sounds are sampled. Either way, the effect is musically distracting, unless you happen to be into exploiting the specific sound modification resources of aliasing. (One person's 'sonic excitement' is another person's 'annoying racket.')

Hiss, rumble, and similar unpitched background sounds that appear not to change as the test tone's pitch rises and falls are the result of either noisy circuits, errors from the limited number of bits per sample, or errors in the timing of when the samples are converted to waveform points. They are generally less objectionable than aliasing, but more objectionable than harmonic distortion. If the loudness is the same for either a strong or a weak sampled sound, then it is certainly more objectionable than would be the case if the background distortion became softer as the desired sampled sound became softer. Test No. 3 sheds light on this aspect of background noise.

Test No. 3. Repeat Test No. 2, but (a) speed up the frequency modulation to about two cycles per second, and (b) put an envelope with a fast attack and slow decay to zero (a couple of seconds) on the test tone. Sample the entire sound. If the entire sound doesn't fit into the memory of the instrument under test, then shorten the test tone's envelope.

Listen to the background noise as the sampled sound drops to zero. The loudness of the noise may go down with the loudness of the test tone. Or the noise may stay at its original volume, and appear to swamp the tail of the test tone's decay. Even worse, the tail may become increasingly crackly and distorted, then suddenly fall into complete silence or otherwise change character abruptly. The first instance is the most desirable behavior, since it is the least noticeable in sound samples that have long attacks or decays. The second instance is just like noise in most analog devices. It's certainly not desirable, but it can usually be managed. The third instance arises from a combination of too few bits per sample and overly simple circuitry that inadequately handles low-level portions of sampled sounds. In any sound with slowly varying amplitude, this last behavior is definitely musically distracting.

Next month, I'll give you some more simple listening tests for sampling instruments. ■