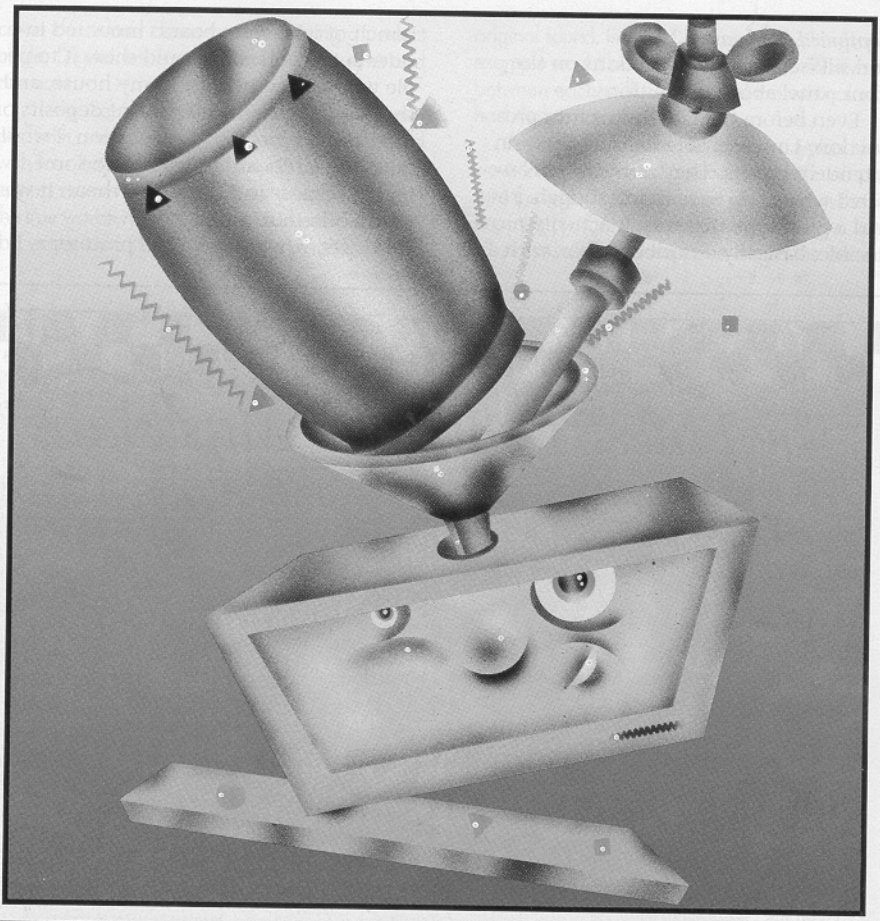

The Art OF Sampling

Part 1—How It All Works

BY NORMAN WEINBERG

WITHOUT A DOUBT, SAMPLERS HAVE BECOME the mainstay in an electronic drummer's arsenal of sound generators. They have been embraced by the drumming community for many reasons—mostly due to their realistic sound. If you've ever used a drum machine, you've worked with sampled sounds—most drum machines play back samples of real drum sounds. But, if you want to create your own drum and percussion sounds, a sampler is the way to go.

Creating realistic drum sounds with a sampler is child's play when compared to synthesizing similar sounds from scratch on any brand of synthesizer. While most synths are competent at creating melodic timbres, only a sampler can produce a convincing snare drum, crash cymbal, tambourine, or a pair of hi-hats gradually opening up after a stroke.

In this three-part series, we'll cover the concepts that allow samplers to perform their feats of magic, we'll take a look at how you can improve your own sampling session, and we'll talk about some creative options for the samples you've generated. All ready? Let's get busy!

Catch A Wave. The act of sampling involves digitizing and recording sound waves. When you kick your bass drum with full force, the sound you're creating is actually nothing more than a movement of pressure waves through the air. Air pressure waves are very similar to water waves.

When a stone is dropped into a small pool of water, waves or ripples are created as the stone "pushes" against the water. As the wave oscillates through the pool, the ripples expand outward. If you look across the top of the water, you'll see that certain spots have a higher level than others.

When a sound is created, similar ripples are sent out through the air, representing higher or lower air pressure. Here's an experiment you can try at home. Take the grill off one of your stereo speakers and watch the speaker's cone moving in and out

Continued on page 44

Continued from page 42

and silk-screened the functions on the front panel above the buttons."

Even before the LM-1 went into production, Linn was able to drum up (pun intended) buyers. They must have wondered what they were in for, though. "I had a prototype that wasn't actually producible, basically a cardboard box with a bunch of wire-wrap boards mounted inside. But it worked. I would show it to people who had come over to my house, and they would give me 50 percent deposits on the finished product. On occasion, I would take this cardboard box down to somebody's session and show it to them. It was pretty hilarious.

"Later, when I had a *real* prototype, I'd keep it in my car. At one party, I showed it to some members of Fleetwood Mac, and I generated some sales from that."

Before the LM-1 ceased production, Linn designed a new drum machine—the LinnDrum—which had many of the LM-1's features but at a significantly lower price: $2,995. "We were horribly afraid that, when we brought out the LinnDrum, nobody would buy the LM-1. So we tried to keep it secret. The funny thing about it is there were a lot of guys who wanted the LM-1. I always thought, 'God, who would want the older technology when the new one's cheaper and better?' But a lot of people liked certain things about the LM-1, like the individual voice tunability."

Linn says there were a total of 525 LM-1s and estimates production of about 5,000 LinnDrums. In 1984, he introduced the forward-looking but bug-laden Linn 9000 sequencer/drum machine. Unfortunately, he wasn't able to remove all of its bugs before Linn Electronics closed its doors in early '86. **D**

Reprinted from Keyboard *magazine.*

Sampling Continued

along with the music. The speaker actually is pushing the air back and forth, creating differing levels of pressure. If you stand close enough to the speaker, you'll be able to feel the pressure waves as they push against your body.

Since you can see and feel these changes in air pressure, there must be a way to measure them. Let's go back to our water wave analogy. One method of measuring this wave would be first to determine the level of the water when the pool is still. Then, you could take measurements of the water level at various times. As a wave crest reaches the measuring stick, there would be more water than the nominal level, and a positive number could be assigned. As a wave trough reaches the stick, there would be less water, and a negative number could be assigned. If it's possible to make this type of measurement with water waves, it should be possible with sound waves. And it *is* possible—it's called sampling.

Analog vs. Digital. In the example above, the water wave itself is analog, but the measurements that we take of the wave are digital. In an analog system, values are infinitely variable. In a digital system, values are finite. The water wave has an infinite number of possible heights, but the measuring stick has a finite number of marked values.

Let's consider a musical example of analog and digital values: A violin uses an analog fingerboard, as the performer can stop the string in an infinite number of places between the nut and the bridge. If the pitch C# is a little flat, the performer can move

that finger up a small amount. A piano, however, is a digital system, as there are no notes available between *C* and *C#*.

Going back to our water wave analogy, we had three pieces of information: the position of measurement (where we put the ruler), the level (the point on the ruler that measures the amount of water), and time (how often we measure the wave). In digital sound sampling, the position is the placement of the microphone, level is the sample resolution, and time is the sample rate.

Sample Rate. Every sampler includes a high-speed clock, and the sample rate is an indication of how often the incoming analog signal is measured. The higher the sample rate, the more closely the sampled sound resembles the original analog sound. For most music systems, the slowest usable sample rate is 10,000 samples per second.

As a rule, faster sample rates produce more accurate high-frequency information. The digital music information contained on a compact disk is sampled at a rate of 44,100 samples per second, and the new R-DAT recorders sample at a rate of 48,000 times per second. In today's consumer samplers, it's not uncommon to find sample rates between 28,000 and 50,000 samples per second.

Sample Resolution. All samplers use a little computer to measure the level of the incoming signal. Since computers use binary bits to represent numbers, the maximum number of usable bits is going to determine how many different numbers are available to measure the analog voltage.

An eight-bit sampler can measure 256 different levels. This means that the analog voltages (in reality, a continuous scale of infinite values) must be assigned one of 256 different numbers. If the true voltage falls between two numbers, it must be rounded off. If a sampler has more values available, less rounding off is required. For those of you into digital trivia, a 12-bit sampler can assign one of 4,096 different values, while a 16-bit machine can assign one of 65,536 values.

As a general guideline, each additional bit of information adds approximately 6dB to the signal-to-noise ratio. This means that a 12-bit machine sounds about 24dB cleaner than an eight-bit machine. It also means that your compact disks (created with 16-bit resolution) have a signal-to-noise ratio of about 96dB.

Visual Examples. In the first illustration, we can see how closely digital data can approximate analog information. All three figures are made up entirely of straight lines. In Fig. 1a [see page 46], nine rectangles are rotated around a single point. Notice how the inner and outer portions of this figure begin to look like a circle. In Fig. 1b, there are 18 rectangles, and Fig. 1c was created with 36 rectangles. As the rate (the number of rectangles) and the resolution (in this case, the accuracy of the printer) increases, the inner and outer portions of the figures come closer and closer to representing a perfect circle. Digital information can fool our eyes and our ears.

In Fig. 2a, an analog waveform is being sampled with a rate of about 2,000 samples per second with four-bit resolution. It's easy to see that the resulting measurements don't accurately represent the true waveform. Fig. 2b comes closer to the analog wave, but still can't catch every nuance of the original sound. Imagine, however, how this graphic example would look if there were 30 to 50 divisions between each millisecond, and the vertical scale had thousands of values.

Putting It All Together. Let's follow a typical sound wave from the original analog changes in air pressure through digitization and playback. In the first step, the cymbal vibrations create changes in the air pressure. The microphone senses the pressure

changes and converts these variations into analog voltages. Next, the analog voltages are measured by the sampler's ADC (Analog-to-Digital Converter), and digitized into binary numbers. Next the bits and bytes representing the waveform are stored in the sampler's RAM (Random Access Memory).

When the sample is played back, the digitized waveform passes through the sampler's DAC (Digital-to-Analog Converter), where each number is transformed back into an analog voltage. The signal then leaves the sampler and moves on to a mixer, amplifier, and its final destination—the speakers.

1a

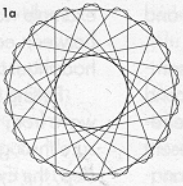

1b

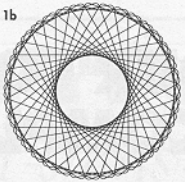

1c

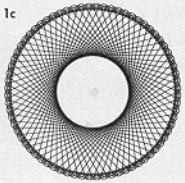

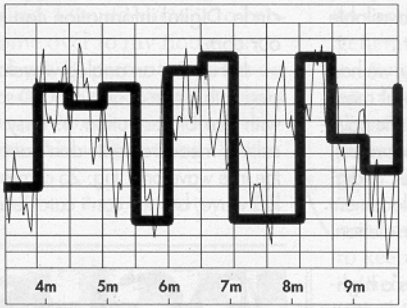

Fig. 2a

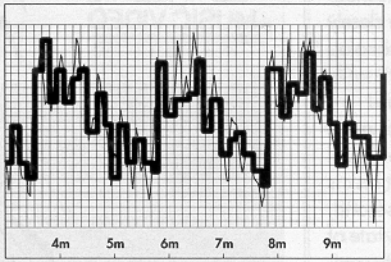

Fig. 2b

Why convert the analog wave into numbers if the numbers are going to be converted back into analog voltages? Aren't we just introducing an extra unnecessary step into the sound production process? The answer, of course, is yes, we are. But, there are several advantages to using digitized waveforms over analog waveforms. Once a sound wave is in digital form, the power of computers can be called into play. And guess what? This is our subject for next month. Stay tuned, samples at 11:00. **D**