REAL-SOUNDS SAMPLING INSTRUMENTS

IN MY OCT. '83 COLUMN, I discussed some of the technical factors behind 'real sounds' machines — instruments that play digital recordings of previously existing sounds rather than generating new waveforms. In this column I will go into a couple of relationships between the 'behind-the-panel' hardware and the quality of the resultant sound. But first, just a little sermon:

A decade or so ago, when real-sounds instruments were mechanical behemoths with tape loops or spinning disks, many people in our industry believed that it was only a matter of time before synthesizers became versatile and sophisticated enough to 'duplicate' acoustic tone sources, at which time real sounds instruments would slip into oblivion. I can remember that, back then, Tom Rhea (formerly a *Keyboard* columnist and currently marketing manager of Moog Music) never subscribed to that theory. He frequently pointed out that a real-sounds instrument is a musical entity unto itself, somewhere between the acoustic 'real thing,' on one hand, and an electronic synthesizer where you shape parts of a sound, on the other. Tom predicted that, one day when technology made it possible for real-sounds instruments to produce high-quality tones at moderate cost, these instruments would enjoy an enduring popularity.

As I've watched musicians getting into the various drum machines, the Fairlight, Bob Easton's 360 Systems keyboard, the E-mu Emulator, and the Kurzweil over the past several years, I've often recalled Tom's predictions. He was (and probably still is) absolutely right. Real-sounds instruments are developing into a class of their own. The best of them combine rich, musically complex timbres with versatile, convenient performance control. Of course, you can't ever sound like Maynard Ferguson from a keyboard, or like Billy Cobham from a bunch of buttons. (This might be the place to point out the obvious: The reasons that you can't sound like them are that (a) you're not them, and (b) what you're playing is not what they're playing.) But you certainly can make good, original, creative music! Using real-sounds instruments to impersonate (but not duplicate) orchestral instruments is just the beginning. End of sermon.

Specs & Sound Quality. Most of us now know that current technology is capable of breathtakingly accurate sound reproduction. Digital audio disks and cassettes store audio with flat frequency response throughout the range of human hearing, and with virtually inaudible noise and distortion. What's especially remarkable is that this technology is readily available to all of us who have a few weeks' pay in our pockets.

Why, then, is it such a big deal to have electronic musical instruments that produce musical tones from digital recordings? Because in musical instruments, all tones have to be instantly accessible. This consideration rules out disks and tapes as the primary storage medium, and leaves us with semiconductor memory chips (called ROMs). A digitally recorded waveform can be pulled out of a ROM quickly enough to satisfy any musician. The drawback of ROM sound storage is that just the ROM chips cost at least several dollars per second of stored sound. In an instrument that offers a choice of timbres and several complete recorded sounds for each timbre, the cost of just the ROMs may approach the kilobuck range.

In designing real-sounds instruments, engineers spend lots of time and effort figuring ways to reduce the amount of memory. Some of these ways involve tradeoffs, while others use tricks of various types. Here are a couple of the variables that instrument designers work with.

Frequency Bandwidth. The audio frequency bandwidth of a digitally sampled sound is determined by the sampling rate — that is, the number of times per second that the audio waveform is sampled. Its amplitude at that instant is converted to a digital number. In the recent draft of the Audio Engineering Society's "Recommended Practice For Professional Digital Audio Applications — Preferred Sampling Frequencies," the sampling rate of 48kHz is proposed as the primary standard for professional audio work. A 48kHz sampling rate is capable of recording an audio frequency band at least 20kHz wide, and therefore imposes no limitation on the system's frequency response. If money were no object, we would like to have real-sounds instruments that have a 20kHz bandwidth. However, in such an instrument, we would need as much semiconductor memory for every second of recorded sound as there is in all of a typical personal computer. An ideal instrument with a stored library of a dozen or so sets of sounds would thus need at least one megabyte of ROM — a huge army of ROM chips.

Fortunately, most musical sounds do not require the full audio bandwidth to sound rich and interesting. In fact, bass and mid-range notes generally sound better in an ensemble if their frequency response is limited. Even high notes often sound better with less than 20kHz bandwidth, if they are equalized with care. Brass and woodwind sounds have very little sound energy above, say, 12kHz, and ensemble strings, believe it or not, do not suffer that much if their bandwidth is limited to 10kHz. Real-sounds instrument design is developing into a serious craft, in which the recording and subsequent processing of the sounds is a delicate, painstaking activity. Tom Oberheim, for instance, recently told me that his company spent nearly two years polishing and refining the sounds of the DX drum machine. Thus, real-sounds instrument designers limit the frequency bandwidth of the sound not only to save memory but also to improve the spectral distribution of the sounds.

Bits Per Sample. Numbers are stored in digital memories as groups of binary (two-state) digits, called 'bits.' A bit is to information processing as an atom is to material processing. Each bit in a sample accounts for a factor of two in the accuracy of a digital number. Thus, if a sample is recorded as an eight-bit number, that means that the dynamic range of the input waveform is divided into 256 (two raised to the eighth power) possible levels, and one of those levels is what is recorded. Of course, audio waveforms vary continuously, so the process of dividing the sampled waveform into 256 levels introduces error. We hear this error as noise (if the sampling rate is not related to the audio frequency) or as distortion (if the sampling rate is some whole-number multiple of the audio frequency). If you follow the math through, you see that each additional bit per sample reduces the error by 6dB. For an eight-bit sample, the error is 48dB below the maximum signal. What does 48dB sound like? It can be as bad as a $20 cassette recorder, recording on discount-store bargain tape. Or, if the audio signal is conditioned properly, if the sampling is synchronous with the audio, and if there is filtering after the sampling, it can sound downright musical. Or anything in between.

Nowadays, most (but not all) real-sounds instruments use eight-bit sampling, and depend on the intelligence and discretion of the instrument designer to sound good. Some top-of-the-line instruments use ten bits, and there are designs on the drawing boards for small instruments that use four bits per sample in combination with some tricky waveform encoding.

As usual in this business, specs are beside the point. The sampling rate and number of bits per sample are variables that instrument designers deal with, not characteristics that directly relate to the sound quality. It is helpful to know what these variables are for a given instrument, because it helps you understand how an instrument works. It's like knowing that a trumpet is made of brass instead of aluminum, because brass has superior acoustic properties and can be worked more easily. In the final analysis, you have to know how the trumpet plays — that is, how good a job the instrument builder did with a common industrial material. Similarly, ROMs, microprocessors, and DACs are the basic materials of our '80s technology. As such, they can be used to make a wide variety of new musical instruments, some of which use stored waveforms. The instrument maker determines how these instruments will sound.

UPCOMING FEATURES:

IN-DEPTH PROFILES OF TWO PROGRESSIVE ROCK GIANTS — TONY BANKS OF GENESIS & TONY KAYE OF YES