ROLAND D-70

MULTITIMBRAL DIGITAL SYNTHESIZER

By Jim Aikin

KEYBOARD REPORT

THE LONG-AWAITED SUCCESSOR to Roland's ground-breaking D-50 has finally hit the streets. When the D-50 debuted in 1987, its sampled attack transients and built-in reverb revolutionized the way we thought about synthesizer sound. Since then, everybody and his brother Hartley has been building synthesizers that sound more or less like the D-50. So is the D-70 another great leap forward in terms of sound power? Is it going to force the rest of the industry to scramble to catch up? Probably not. What it sounds like, basically, is another D-50. Well, no—it sounds like what we *remember* the D-50 sounding like, only richer and fuller and more contemporary than the real thing, and with less inherent noise.

Nothing succeeds like success. Apparently, Roland's goal was to keep the classic sound intact while adding a whole lot of features that the D-50 didn't have—multitimbral operation, extra expression sliders, four-channel MIDI transmission, new waveforms, multimode filters, filters on the tones that use sampled waveforms, more flexible programming architecture, extra audio outputs, a 76-note keyboard, and a bigger LCD window. Lots of mouth-watering treats, especially for the advanced user. There's also a new type of waveform generator, which we'll have more to say about below. What the D-50 has that the D-70 doesn't are separate chorusing for upper and lower tones, built-in programmable EQ, and several obscure but sometimes useful programming features.

Overview. The almost-highest level of the D-70's architecture is a set of 64 "performances." (Watch out, here comes Rolandspeak.) Each performance is six-channel multitimbral, with a separate patch (what you and I would call a patch) for each of five incoming MIDI channels and a sixth rhythm channel to make the drum noises. When performing, you can either select one of the performances, or bypass this level and select one of the patches directly. For ordinary use, the latter may be preferable, as there are 128 patches in memory. The downside is that when you select patches directly, you'll always be working with the same effects sound, because effects are programmed at the performance level, not at the patch level.

Each patch is made up of four "tones." As in the D-50, each tone has its own parameter settings for envelopes, LFO effects, keyboard scaling, and so on. Two differences immediately become apparent, however: First, all tones have active multimode TVFs (time variant filters), so you can filter those PCM spectrum waves to get a better blend. This is definitely an improvement. And second, instead of being stored with the patches, tones are independent entities. The machine stores 128 tones, and a patch accesses four of the existing set.

This is both good news and bad news. It's good news because if you program a tone that you like, you can easily insert it into several patches without having to copy the parameter values. And you can very easily invent new patches simply by combining existing tones. It's bad news because, if you program a new patch from scratch, developing four new tones that blend well with one another, where are you going to store the tones? The D-70 has no utility command for listing which of the stored tones are used in which patches; thus there is no way to tell, during a programming session, which tones are unused, or used only in patches you don't care about, and thus can safely be overwritten. When you store your new patch, you may be drastically altering the sound of dozens of other patches! The only way to avoid this is to use pencil and paper to keep a cross-indexed log listing which of your patches use which tones. (Cross-indexing will be coming soon, we're sure, to an editor/librarian near you.)

If we ignore this major oversight, the architecture makes good sense. What's especially cool is that the patch parameters include offsets for some of the more commonly used parameters of the tones: Loudness, panning, tuning (coarse and fine), filter cutoff and resonance, and envelope attack and release can be individually offset for each tone in a patch. Thus a single tone could be slow, loud, and muted in one patch, and fast, quiet, and bright in another. In some situations, this may allow you to program a patch precisely to your taste without having to touch the tone parameters, thus avoiding potential foul-ups.

Above the level of performances is a set of ten "user sets." A user set contains nothing but a list of five performances, which can be accessed in play mode with the five function buttons under the display. The names of the performances are displayed as a menu in the window, making it very easy to see what you're doing.

A set of four data sliders can be used for many types of editing at the patch level, though not when you delve into editing tones. These sliders are echoed in the display window with a little graphic that shows you both their current physical position and the current value of the parameters that they are assigned to. The slider will be inactive until the physical position crosses the value in memory; at that point, the two start to move together. In addition to editing, these sliders can be used very effectively as expression controllers in live performance, as they will adjust the filter cutoff, resonance, fine-tuning, or loudness (though not the coarse tuning or pan position) of a sustaining note. It's too bad you can't set up your own programmed bank of functions for the four sliders; as it is, they will all four perform the same function, with one assigned to each of the tones. They will transmit and respond to MIDI controllers 16 through 19 on a special "control channel," which is great, as it lets you automate the balance of the four tones during a sequence. A few extra keystrokes are required to set this up properly, but for the few, the proud, the MIDIots, it will be worth the trouble.

The D-70 can have up to 30-voice polyphony, but the actual number of notes it will play at a time depends on how many tones (oscillators) are active in the various patches. If you're playing a three-tone patch, you've got ten-voice polyphony. With a full four tones in a patch, the polyphony is actually one note

ROLAND D-70

Description: Multitimbral synthesizer.
Keyboard: 76 notes, *E* to *G*. Velocity, release velocity, and channel pressure. 30-tone polyphony (actual number of notes depends on how many tones are active in the patch). Two left-hand expression sliders in addition to standard Roland pitch-bend/mod lever.
Memory: Ten user sets (five performances each), 64 performances (five-part multitimbral plus drum part), 128 patches, 128 independent tones. One drum kit. 3Mb waveform ROM. Reads U-110 series PCM cards.
Voice Architecture: Up to four tones per patch, each tone with one oscillator, resonant multimode filter (low/high/bandpass), two four-segment envelopes, attack/release pitch envelope, and one LFO. Offsets in patch memory for some tone parameters.
Features: Built-in programmable reverb and chorus. Six-channel multitimbral operation. DLM wave generator. Independent assignment of performance controllers to internal voice response and MIDI data transmission. Real-time slider control of filter and loudness of sustaining tones. Four-channel, four-zone MIDI transmission. Graphic display of envelopes.
Interfacing: Two PCM card slots, one RAM card slot. Headphone out, four audio outs (two stereo pairs, one processed by effects, one dry). Sustain pedal, footswitch, and expression pedal inputs. Rear-panel LCD contrast knob and memory-protect switch. MIDI in/out/thru.
Dimensions: 46-3/4" x 12-1/4" x 3-1/4". 26.5 lbs.
Suggested Retail Price: $2,395.00.
Contact: RolandCorp US, 7200 Dominion Circle, Los Angeles, CA 90040-3647. (213) 685-5141.

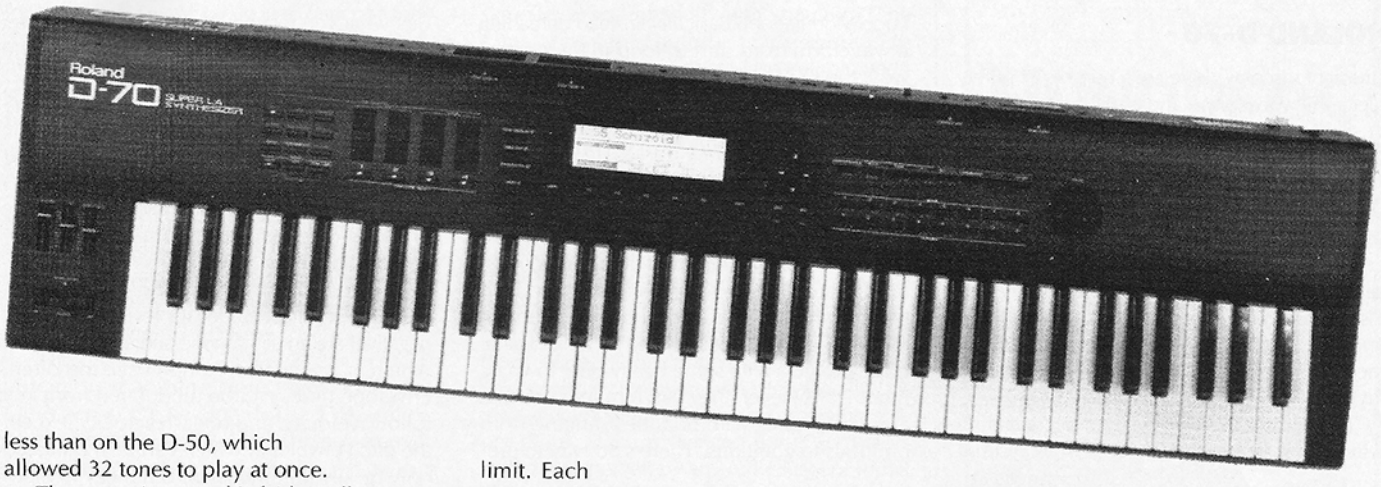

less than on the D-50, which allowed 32 tones to play at once.

The owner's manual is fairly well written, but perhaps not as well organized as it might be. Frequent cross-references to various other sections help you track down related information that's scattered in different sections. The operating system itself isn't exactly linear, either. Buttons change their functions or become inoperative depending on what page you started from; some data entry fields are redundant with other fields on other pages, yet in some pages you can't access data that is displayed; it's there for reference purposes only. On the other hand, there are lots of convenience features, such as graphic display of envelopes and the ability to give each MIDI channel a name (to show at a glance which external module you'll be playing). Each window 'remembers' where you left the cursor, which is great. Not so great: The only way to switch from one tone to another within a patch when editing tone parameters is to scoot the cursor up to the field that shows the tone number and then use the inc/dec buttons. There are tone mute buttons on the front panel, but no tone select buttons. What it boils down to is that this is a complicated piece of equipment, one for which the learning curve is not exactly a bunny slope. Experienced synth programmers shouldn't have much trouble mastering it.

Performance Programming. Each performance contains six "parts." The first five are assigned patches, while the sixth is the rhythm part. Each performance also stores certain other parameters, such as offsets (discussed above), the effects, and a patch parameter called "analog feel." Analog feel seems to be merely a pitch randomization function that is applied separately to each tone. (The Yamaha DX7II's random pitch function, while similar in concept, affected all six operators together. We like Roland's version better.) The result is that both the center pitch and the amount of detuning change from note to note. Small amounts of analog feel seem to impart a pleasing liveliness to the sound; large amounts are merely silly.

You can give each part a keyboard zone, and program whether it will recognize or ignore incoming MIDI volume and program change data. This could be very helpful for complex sequences, as it lets you cheat the 16-channel limit. Each part can also be sent through the effects section, or bypass it.

We were delighted to discover that the five patches assigned to a performance can all be edited in a single session; no need to save one before you move to the next, as there are five edit buffers. This is one of those places where powerful software can confuse the unwary, however: If the same tone is used in several of those patches, you can edit one and store the edit without hearing changes in the others—until the next time you call up the performance from memory, which will cause all of the edit buffers to be reloaded.

PROS & CONS

Pros: Four-channel, four-zone MIDI transmission. 76-note keyboard. Clean sound. Large graphic display. Real-time filter cutoff control from front-panel sliders.

Cons: Sluggish response to heavy MIDI data. LFO rate slows down in dense passages. No amplitude scaling of tones across keyboard. No edit compare function.

Patch Programming. Most of the parameters in this section will be familiar to Roland users. There are a few innovations, however, and at least one curious omission. As already mentioned, a few of the more often used parameters are programmed at both the patch level and individually in the separate tones. These include filter cutoff and resonance, loudness, envelope attack and release, and coarse and fine tuning. These parameters can be offset separately for each tone in a patch. Because of this, if you're programming a tone from scratch we'd recommend that you zero out all of the values, so as to avoid surprises when you later try to use the tone in another patch. A patch also contains global settings for overall velocity sensitivity and velocity curve. Since the differences among the velocity curves are fairly subtle, we aren't real sure what happens when the patch is programmed with a different curve than is called for by the individual tones.

The tones can be programmed in five basic areas: waveform, pitch, filter, amplitude, and LFO. Let's start with the waveforms, since Roland has added a new feature here that programmers may want to know about.

Instead of the attack transients (with the fundamental filtered out) that were the norm in the D-50, the D-70 has a variety of full looped waveforms. Most of the loops are quite good, as is the matching of multisamples across the keyboard. (The D-50 didn't have multisamples at all.) A number of waves don't play in the top octave—in a few cases, the top two octaves. Okay, so they're working with a fixed amount of memory for waves. If we had to choose between having more types of waves with silent octaves and fewer types of waves with more multisamples to cover the high register, maybe we'd make the same choices Roland did. But it does seem a little strange to build a 76-note keyboard that turns into a 58-note keyboard when you program an organ sound. You can transpose the whole keyboard down an octave or two to get the upper keys back, but then you get the famous rubber trombone effect in the bottom octave.

Speaking of organs, they got around the fact that the D-70 has no Leslie simulation in its effects section by multisampling an organ with the rotors cranking. It sounds pretty ballsy, but of course you have no control over the speed or depth of the Leslie effect, and the speed changes as you play from one key to another. Likewise, there's no distortion effect, but there is a very good fuzz guitar sample.

The basses are very good. The lower end of the choir, unfortunately, seems to be the same horrible pitch-wobbling sample that was used in the U-110. Can we just quietly take this one out behind the barn and shoot it, guys? Several of the waveform options, such as the slap basses, offer alternate keyboard layouts (with different split points) of mostly the same group of samples. Curiously, one pair (strings 1 and strings 2) seemed to be entirely identical from one end of the keyboard to the other. The flute was sampled with vibrato, which you may or may not like. We're happy to report that, unlike the D-50, the D-70 has quite realistic piano samples, both acoustic and electric.

In addition to the acoustic and electric in-

strument samples, there are a number of purely electronic waveforms, and also 28 different percussion sounds. The latter can be used either in drum kits or as attack transients. There are five snares, four kicks, three toms, the TR-808 snare, hi-hat, and "tom," and the usual assortment of other stuff. Missing: Timbales, bongos, and congas. Fortunately, the U-110 PCM card series has a Latin percussion card. These cards also have fixed drum kit layouts, as we recall, and we didn't get a chance to test exactly how the samples will be integrated into the keyboard layout of a D-70 kit.

Now, about this new method of generating waveforms. Roland calls it DLM (Differential Loop Modulation). Basically, instead of reading a waveform from start to loop in the normal way, the D-70 can start at an arbitrary point in the wavetable and read a segment of arbitrary length. Instead of looping this directly, however, which might result in a high-pitched, low-amplitude sound, the instrument intelligently repeats the selected wavedata in a staircase. This results in a longer wave loop (in some cases as much as two or three seconds) that contains assorted high harmonics. Without filtering, these waves all sound buzzy and unpleasant—maybe just the thing for an industrial noise project. With some heavy-duty lowpass filtering, however, they can turn into unusual animated tones with qualities ranging from ominous to gorgeous. There's no way to predict what you're going to get with this feature: just fiddle with the parameters until you find something that you like.

The filter and amplitude envelopes are actually a little less flexible than those on the D-50, with four rates and three levels rather than five of each. The pitch envelope offers only attack and release rates and times in place of the D-50's four-rate, four-level pitch envelopes, but each tone has its own pitch envelope, which is not the case on the D-50. Also, the smallest values for level (positive or negative) are an eighth-tone and a quarter-tone, which is a subtlety you don't see too often. Envelope times can be modulated from key follow, velocity, and release velocity (except the pitch envelope, which can be modulated only by key follow). There is no way to modulate one segment of an envelope without modulating them all, except that release velocity always modulates release time. What's worse, the modulation of envelope times is not bipolar. High velocities always make the envelope shorter, which is great for attack times, but fairly silly for decay times, where you would prefer to have hard velocities create longer notes rather than shorter ones. Likewise, release velocity modulation always results in quicker releases creating quicker release times. This is normally the polarity you'd want, but the effect can't be inverted.

The longest LFO time is in excess of five minutes (!) for one complete cycle, but normal vibrato lies in the range between 75 and 108 (full parameter range 0-127). There is no choice of retriggering options in the LFO section, but both delay and rise time are programmable. The modulation amount can be positive or negative, and modulation can be applied in individual amounts for pitch, filter, and amplitude. Each tone has only one LFO, however, which means you can't program separate mod rates for the filter, amplitude, and pitch of a single tone, as on the D-50. Only four LFO waveforms are offered, but an offset parameter lets you program square-wave trills that go only sharp or flat from center pitch.

We only discovered the big problem with the LFOs by accident. We set up a two-tone patch that used some heavy wah-wah-type LFO modulation of the filters, and played some repeated four-note chords. In this situation, the LFO rate slowed down to less than half of what it was supposed to be, presumably because the processor was working so hard calculating the interactions of envelopes with LFOs. This was not a torture-test, either; it was a very normal musical application, and the D-70 couldn't handle it.

The D-70's velocity curves are not the most comprehensive we've seen. There are four curves—linear, two concave (suitable for players who play hard), and an S-shaped curve that offers the greatest sensitivity in the middle of the velocity range. There are no convex curves. Velocity sensitivity can't be inverted, so it's not possible to do the standard type of velocity cross-fades. Velocity cross-switching is supported, however. The D-70 also offers a more interesting type of velocity response called velocity mixing. This allows you to hear one tone at low velocities and have this tone be dominated by a second tone at high velocities, while both are responding to high velocities by getting louder.

The same type of fractional pitch scaling found in the D-50 has been retained, but redesigned slightly. There are fewer options between 0 and 100%, except that a couple of slightly compressed anti-stretch tuning curves have been added, as have a couple of reverse scales. The filter can track the keyboard in adjustable amounts, of course.

The most significant weak link in D-70 voicing is that there is no level scaling of tones across the keyboard. The TVA scaling parameters are applied strictly to the envelope times, not to the levels. What this means is that the D-70 offers no easy way to bring a component of the sound in or out smoothly at one end of the keyboard. You can use filter keyboard tracking, but this may not always produce the sound that you want, as it will remove the highs (or lows) instead of dropping the whole amplitude uniformly. Even the filter tracking doesn't have the D-50's break point/level scaling, only a uniform (though programmable) keyboard tracking amount. It's a mystery to us why they would have lopped off such an important parameter; maybe they felt that since a tone might be used in many different patches, programmers wouldn't often want to be tied to the same level scaling in all of the patches. (The same argument could be made about LFO depth, of course, and that wasn't eliminated.)

Percussion Section. The D-70 contains only one global percussion layout, but the features for programming it are fairly extensive. You can assign a separate drum tone to each key (no layering of drums, sorry) and tune it up or down to the transposition limits of the sample—several octaves each way. Fine tuning has a very high resolution. Non-percussive samples can also be made part of the kit, if you like. Each key can be given its own filter settings (including filter mode, cutoff, and resonance) and envelopes (filter, amplitude, and pitch attack). Panning and output routing are also programmable per drum, and any key can be set so that its sound is turned off by any other key (for hi-hat closures). LFO effects are not available for the drums, however.

Built-In Effects. Not the best effects section we've ever seen in an over-$2,000 instrument. Don't get us wrong—it will work perfectly well to enhance the sound of the instrument. But it doesn't have niceties like high damping in the reverb, there are no algorithms for distortion, Leslie simulation, or left/right delays with separate delay times, and the maximum delay time is not very long. What you've got here is your basic reverb and chorus, which can be strung together in series or in parallel, allowing you to send one of the parts in a performance through each of them, or combine them for a richer sound. (Parts can also be routed dry or to the direct outputs, which are also dry.)

The reverb section includes three rooms and two halls, plus gated reverb and two delay algorithms. Parameters include output level of the effect (which has no effect on the dry signal level—this is not a wet/dry mix control), time, and feedback for the delay. Maximum delay time is under half a second, and delay time is adjustable not in milliseconds but in arbitrary values from 0 to 31. For the chorus section, you get your choice of chorus 1 and 2 algorithms, plus chorusing with feedback, flanging, and a short delay.

MIDI Implementation. As a master keyboard, the D-70 has a lot to offer. The keyboard can be split into four zones, each with its own channel, transposition, and values for volume and program number. Each performance can be programmed with its own MIDI output parameters, and the volume and program number will be transmitted when the performance is called up from memory. The MIDI output can also be assigned a velocity curve, velocity sensitivity, and a choice of velocity response modes (cross-switching or cross-mixing, with a programmable threshold value) for the upper and lower halves of a split.

Controllers can be remapped internally (modulation to aftertouch, aftertouch to pitchbend, etc.), and can also be assigned to transmit other types of MIDI data. Because these two functions are independent, you can easily set up a situation in which the data recorded into

your sequencer will sound different on playback, but this isn't a problem, it's a feature. You just have to know how your performance presets are programmed.

With five different patches available in a single performance, and each patch assigned to its own MIDI reception channel, you might expect to be able to play four patches at a time from the keyboard by using the MIDI output zones. Alas, there is no way to do this directly. Only one patch can be assigned as the "keyboard part." In order to get four-way splits and layers of whole patches, you have to run a MIDI cable from the instrument's output back to its input. (You can get four-way splits within a single patch, but these will contain only a single tone (oscillator) per note.)

We don't recommend doing four-patch layers, as the instrument gets abominably sluggish when you set it up this way. How sluggish does it get? We couldn't get any hard data by hooking the synth up to a sequencer, because the MIDI implementation has been set up so that all of the keyboard notes are transmitted first, before the processor starts turning on notes. Our sequencer data looked great: Four-note keyboard chords, transmitted on four channels for each note, appeared as perfect 16-note vertical stacks, with zero clocks between note-ons. The sound of the instrument was something else again. This was a torture-test, not something you'd do every day, but with a three-tone patch assigned to four layered channels, four-note keyboard chords (12 oscillators per key) were smeared by what sounded like 200 milliseconds or more, and if we banged out some fast chords with this layout, the D-70 continued to turn notes on for a quarter-second or so after we stopped playing. Conclusion: It's fast enough for simple multitimbral chores, but don't ask it to run the hundred meters against a photon.

We had some trouble at first figuring out how to do basic multi-channel recording into a sequencer from the D-70, because of the large assortment of different screen pages that allow you to do different things. Programming your five parts to receive on separate channels and then choosing a different part to be the "keyboard part," *i.e.*, the part played from the keyboard, will *not* change the MIDI transmit channel. The way to use a sequencer turns out to be to use the sequencer's software thru and turn the keyboard part to 'off.' At this point, the keyboard will not play the internal voices at all. You can then switch to the MIDI output page and use four of the front-panel buttons to select among the four output zones (which in most cases will be set to separate channels), and when the D-70 sees its own MIDI data coming back from the sequencer, it will play whichever part is assigned to that channel.

This isn't quite the whole story, however. Even when you switch the keyboard part to off, you'll always be transmitting on the global "control channel" in addition to whatever MIDI output channel you select. You can turn the control channel off, but then the D-70 won't switch to new performances when it receives program changes on this channel. A better alternative is to use an undocumented power-up option (switch the unit on while holding down number key 4) to disable the control-channel note transmission.

We discovered a tiny but annoying bug in the MIDI output programming section: We set up a four-channel layer with a transposition value for each layer, and when we changed program numbers, the volumes for the outputs would sometimes change as well. In the absence of the transpositions, the volumes remained unaffected by our other edits. Oops. It's easy enough to reset the values manually before saving the performance to memory, so this should cause no problems.

Conclusions. I wish I could spend two or three months working intensively with the D-70 before writing this paragraph. It's a deep instrument, not one that you can make snap judgments about or characterize accurately in a few quick phrases. But I know most *Keyboard* readers would rather not wait that long for a review, so here goes.

Here's one bunch of quick phrases: This is a fully professional synthesizer. It sounds great, with excellent raw samples and comprehensive patch programming options. Its 76-note keyboard and assignable controllers are complemented by a large, easy-to-read graphic display. The four-channel MIDI output, with 64 programmable configurations, makes it close to ideal as a master controller. The innovative "user set" arrangement is bound to make live performing a breeze. And power sequencer users are sure to appreciate being able to edit five patches at once.

Another bunch of quick phrases: In a fully professional synthesizer, we can't help expecting better programmable effects than the D-70 delivers. The slowness of the response to incoming multi-channel chords is worrisome, and the inability to keep up with its own LFOs in two-tone, four-note chords is appalling. We're more than a little surprised that you can't split and zone patches across the keyboard without using a MIDI cable as a hardware kluge. The operating system is a mixed bag—lots of features, but they're laid out in a rather confusing manner. With 128 independent tones to choose from in building patches, it's infuriating that you can't tell which patches a tone is used in without going through all 128 patches by hand and writing them down. And the drum kit parameters are excellent—so why is there only one global drum kit?

The bottom line? It's a long way from perfect, but it's also the best synthesizer Roland has ever produced. And considering how many great instruments they have in their line, that's saying quite a bit. ■