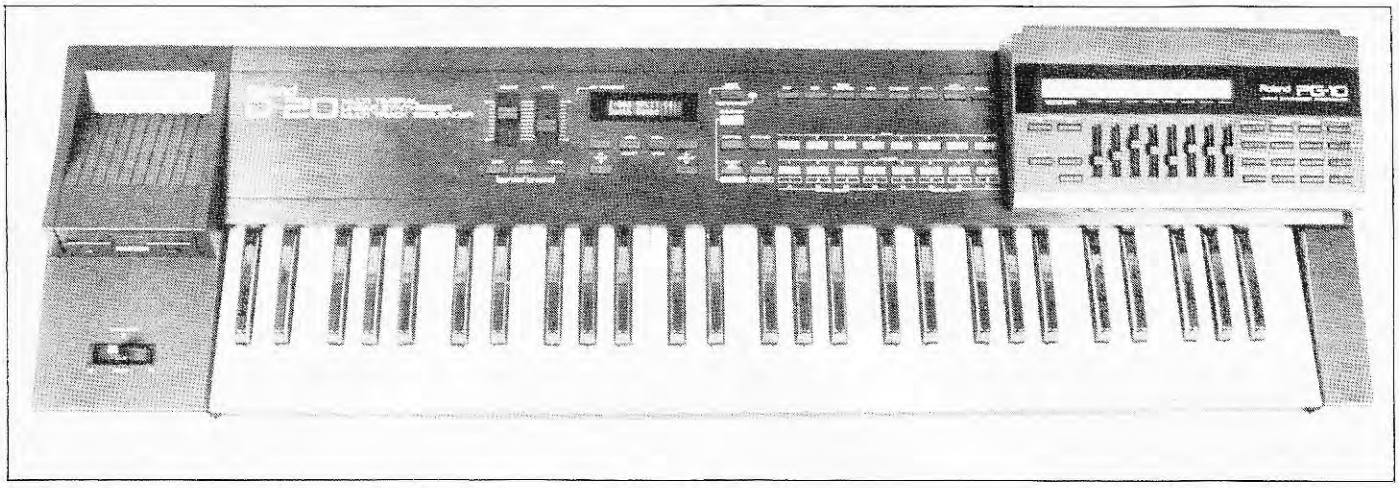

ROLAND D-20 DIGITAL SYNTHESIZER

By Mark Vail

THE WAR OF THE ALL-IN-ONE SYNTH machines is escalating. A typical battlewagon's weaponry includes a five-octave keyboard, multi-timbral sound generation, and a basic MIDI sequencer. Extra firepower—in the form of ROM-based percussion samples, on-board reverb, and a built-in disk drive—is not yet standard issue, but is highly prized in battle. In recent front-line skirmishes, Ensoniq fired off its SQ-80 synth and Korg retaliated with its M1 sample player. Now Roland has joined the fray, wielding its new D-20 synth. Meanwhile, the poor working musician is stranded on the battlefield wondering which camp to join.

In terms of its sound, the D-20 comes from the same lineage as the popular D-50 (reviewed in *Keyboard*, Sept. '87). Roland pared away some of the D-50 features and added others. The result is a $1,795 instrument that may appeal to a different musician than the D-50 does. While the D-50's popularity is rooted in the thousands of performers and recording artists who took to its powerful sound, the D-20 is targeted at the musician looking for a multitimbral instrument with sequencing. (See the sidebar on page 135 for a discussion on the differences among Roland's LA synthesizers.) The question is, how does the D-20 stand up on its own as a do-it-all machine?

Keyboard & Controllers. The D-20's five-octave velocity-sensitive keyboard feels as good as any organ-style synth controller. Sadly lacking is aftertouch sensitivity—okay, we admit that we've been spoiled by all the synths that do respond to aftertouch. The D-20's keyboard can be split at one variable point.

For modulation and pitch-bend control, the D-20 offers the familiar Roland paddle. Moving it left and right bends the pitch down and up, and pushing it forward adds modulation. The paddle's forward throw is very short, but it's pressure-sensitive. If you're used to a Moog-style wheel with lots of range for gradually imparting more modulation, Roland's paddle will take some getting used to.

Voice Architecture. Along with linear arithmetic synthesis technology come unique terms for its building blocks. D-50 partisans won't recognize some of these, because they have been added to describe components in multitimbral operation, one of the D-20's two major modes. In performance mode, the D-20 works much the same way that the D-50 does.

In Roland parlance, a *partial* is the equivalent of a simple synthesizer. A partial either plays a PCM recording or produces a synth waveform. A PCM partial consists of a sound generator and an amplifier. A synth partial consists of a sound generator (which produces either a sawtooth or a variable-width square wave), an amplifier, and a lowpass filter. You can control the synth partial's pulse width dynamically using velocity. The sounds of both the square and the sawtooth waves are affected by pulse width modulation.

Where Roland's D-50 had a healthy supply of three LFOs per voice for modulating pitch, pulse width, filter cutoff frequency, and amplifier level, the D-20 only provides one LFO, and its sole purpose is to add vibrato to a partial. You can only change the LFO's rate, depth, and sensitivity to the modulation paddle. Though the manual doesn't tell you this, you can delay vibrato using the initial time value of the pitch envelope. The remainder of the D-20's partial parameters are identical to those on the D-50. If you want an in-depth description of these details, please refer to our review of the D-50.

At the next hierarchy is the *tone*. The D-20 comes with 128 preset tones, and you can store 64 tones of your own. A tone consists of two pairs of partials and a collection of *common parameters*. One of these parameters, the *structure*, defines the type of sound generator for each partial in a pair, as well as how their signals will be routed. Where the D-50 only features seven different structure types, the D-20 has 13. Two of the new structures allow you to route the two partials' signals to separate sides of the stereo mix; the other four route both signals through the ring modulator.

Next up the line is the *timbre*. A timbre is

Roland D-20

Description: Digital synthesizer/sample player with built-in sequencer, drum machine, and disk drive.
Keyboard: Five octaves, C to C. Velocity sensing. Layer and split modes, with variable split point.
Voices: 16 two-partial voices, eight four-partial voices. Dynamic voice allocation in multitimbral mode.
Memory: 128 patches, 128 timbres, 128 preset tones, 64 programmable tones, 63 preset rhythm sounds. Sequencer: 16,000 notes. Rhythm section: 85-key drum-kit, 32 preset patterns, 32 programmable patterns, 96 notes per pattern. Optional memory card (M-256D): 128 patches, 128 timbres, 64 tones, 32 rhythm patterns, one rhythm track sequence, one rhythm setup.
Interfacing: MIDI in, out, thru. Stereo 1/4" audio outs and headphone out. Three footswitch inputs (sustain, sequencer start/stop, and punch in/out). Memory card slot. PG-10: MIDI in, out, thru, parameter in.
Features: Eight-voice multitimbral operation. Eight-track sequencer. PCM drum sounds and rhythm sequencer. Built-in 3½" disk drive. Digital reverb with eight presets.
Dimensions: 40" x 11⅞" x 4¼", 22 lbs, 4 oz.
List Price: $1,795.00. PG-10: $399.50.
Contact: RolandCorp, 7200 Dominion Circle, Los Angeles, CA 90040. (213) 685-5141.

a single-tone sound module with transposition, bend range, and reverb on/off controls. Roland provides a timbre assignment function so that you can either choose to let voices within a timbre be stolen when they are needed or allow no new notes to play until a voice is available. You can also decide whether a note that you play repeatedly stacks new voices on the note, or merely mutes and retriggers the same voice.

An eight-timbre *setup* is the D-20's highest level in multitimbral mode. Each timbre in a setup is assigned a volume level and stereo pan position. Global parameters for all of the timbres in a setup include the tuning, transposition, and keyboard split point, as well as the reverb type, level, and time. There are eight types of reverb, plus a choice of no reverb for the setup. Reverb types are small and medium rooms, medium and large halls, plate, and three kinds of delay: a single slapback echo, a fading echo, and a fading echo that pans left and right.

You can layer all eight timbres and the rhythm tones in multitimbral mode by setting their MIDI channels to the same channel the keyboard is transmitting on. Since you can set the keyboard to send on different channels above and below the split point, you can assign any timbre's MIDI channel to play on either side of the keyboard.

A sound in the performance mode is called a *patch*, which is made up of two tones that can be layered across the keyboard or assigned to opposite sides of the programmable split point. Alternatively, a patch can use only one tone, which doubles the number of voices available on the keyboard. Each tone in a patch features the same parameters as those in a timbre, along with pitch-bend range and assign mode. In addition, you can adjust the volume level of the two tones in relation to each other.

As with the D-50, when you edit any parameter on the D-20, you must store it before calling up another sound or turning the instrument off. Before you save a tone, timbre, or patch, you can use the compare button to hear the program that you will be replacing. One neat D-50 edit function that didn't make the transition to the D-20 is the undo feature.

Samples. Where many digital synthesizers use a single cycle of a sampled waveform, the D-20 has 144 looped and 112 non-looped PCM recordings. Although there are some samples taken at low and high ends of an instrument's range, there are no multisamples in the D-20's arsenal.

In our usual heartless manner, we carefully scrutinized each of the D-20's 256 PCM sounds. We found that the majority of the samples are noisy, and a good many of them are very noisy. Granted, in these tests we singled out each sample and played it very loud. The best samples included most of the drums, although there were some audible problems in a few of these, most noticeably the conga. The noisiest samples were of the piano, bass, guitar, organ, and voice.

In non-critical live performance, the inherent sample noise probably won't matter. However, in critical recording situations, the D-20 may not sound as clean as some people would prefer.

Rhythm Section. The D-20 comes with 63 preset rhythm tones, and you can also use

Pros & Cons

Pros: Multitimbral LA synthesis with built-in reverb. Built-in sequencer with drum machine programming. Strong drum sounds. Built-in 3½" disk drive.

Cons: One sequence in memory, volatile sequencer memory. Disk drive won't store MIDI sys-ex data for external devices. No aftertouch. Poor manual.

WHAT'S THE THE D-DIFFERENCE

By Mark Vail

TWELVE DONUTS IN THE BOX, SOME chocolate-covered, some filled with custard, some sugar-frosted, each one uniquely enticing. Despite their differences, however, they all share one basic ingredient: sugar.

Roland's LA synths, which we'll refer to as the *D*'s, include the big daddy D-50 keyboard synth, its rack-mount brother D-550, and the baby D's: the D-20 and D-10 keyboard synths and the D-110 rack-mount. Although Roland's MT-32 isn't branded with a D, it is directly related to the family. The common ingredient here is LA synthesis, and we all know how sweet that sounds. But what sets each of these devices apart from the others?

Let's take a closer look, starting with the **D-50**, and work our way down. To refresh your memory, the D-50 ($2,095) is 16-voice with split and layer modes, a five-octave keyboard with velocity and aftertouch response, stereo outputs, and built-in digital EQ, chorusing, and delay effects. It also has a front panel joystick that is useful for blending the partials in a sound. The **D-550** ($1,850) doesn't have the keyboard or joystick, but it uses the same sound generation circuits as the D-50, along with the added benefits of programmable chase play on/off and portamento on/off per patch. The daddy D's have lots of goodies, but they lack one function that many of their competing products feature—neither unit is multitimbral.

Enter the baby D's, who love playing the multitimbral game. Each of them will give you up to eight different sounds at once through their stereo outputs, and voice allocation is dynamic—you don't have to assign a specific number of voices for each MIDI channel. What's more, the baby D's can store 64 programmable sounds to go with their 128 presets, compared to the 64-patch limit of the daddy D's. Finally, the baby D's have 13 different programming structures (a structure being a way of combining and routing partials). The daddy D's offer only seven structures.

What about individual differences among the baby D's? The **D-20** ($1,795) has its own 3½" disk drive, sequencer, and, like the **D-10** ($1,395), a programmable drum machine. The **D-110** ($995) doesn't have a sequencer or a drum machine, but it does feature the same 63 percussion sounds found in the other baby D's. Its MIDI configurations are programmable. The home recordist will also appreciate the D-110's additional six assignable outputs. You can assign any sound to one of the programmable outputs, and it will bypass the internal reverb, allowing you to process it any way you see fit.

The **MT-32** ($695) features a synthesis architecture similar to the one used in the baby D's, but its signal-to-noise ratio isn't quite as good and its drum sounds—of which there are only 30—don't sound as realistic as those in the baby D's. There are 128 presets in the MT-32, but to program new sounds you'll need a software editor, because the MT-32 can't be programmed from its front panel. It also doesn't store any of its performance settings when you shut it off, so you'll either need to reset it manually for each session or program your sequencer to send program changes, MIDI volume settings, and various system-exclusive messages for each new song.

As you can see, you pay less for a baby D or the MT-32 than you would for a daddy D, but you get some extras. So what gets sacrificed for those extras? To start with, the little guys won't respond to aftertouch. The D-50 and D-550 provide three LFOs per voice, but the baby D's and the MT-32 have to depend on one per voice. The D-50's data display is better. And you won't find the EQ and chorusing effects of the daddy D's in the little guys, because the digital reverb used in the baby D's is different. It's only 12-bit. Not only do the EQ and chorusing enhance the sound of the daddy D's, but the signal stays digital all the way through the instrument until the outputs. The signal in a baby D starts out as digital, then gets converted to analog and back to digital when it goes into the reverb. This extra conversion process degrades the sound quality somewhat, as Peter Gotcher explains in his Digital Sampling column on page 125. The MT-32 has the same reverb as the baby D's, but the surrounding circuits aren't of the same quality; therefore, the MT-32 doesn't sound quite as good as the baby D's. As for the number of reverb and delay settings offered by each type of device, the daddy D's give you 32, the baby D's eight, and the MT-32 four.

Can you hear the difference in sound between a baby D and a daddy D? You sure can. The D-50 and D-550 produce sounds that jump out and tickle the eardrums. The baby D's make pretty sounds—they just don't sound as remarkable as the daddy D's.

What about programming? If you own either the **PG-1000** programmer ($495) or a software editor for your daddy D, don't expect to use it on a baby D, because the two groups are incompatible. However, the **PG-10** programmer ($399.50) can be used on any of the baby D's, as can editor software. Unfortunately, the MT-32 requires its own editor program, and it won't work with the PG-10. As for RAM card storage, the daddies and babies use the same kind of card, but the card memory has to be formatted for one or the other type of machine. The MT-32 doesn't have a card slot. ■

any of the 64 programmable tones in programming rhythm patterns. Sixty-three rhythm tones are mapped across the keyboard. You can play these in manual play mode. Since the D-20 only has 61 keys, you'll need to transpose the keyboard up or down a half-step to access the 62nd and 63rd tones. You can assign rhythm tones all the way from *C1* to *C8* by transposing the keyboard up and down and assigning a sound to the keys. You can reprogram the preset drum sound layout very easily. When you edit the key assignments, the D-20's display shows the note that you play (by name and octave), the tone that it controls, the volume level and pan position of the tone, and whether reverb is on or off for the tone. Any of these parameters is accessible for editing by moving the cursor to the parameter. The cursor will stay on that parameter for any new key that you access. Our experimental juices started flowing when we found we could play a rhythm pattern on the keyboard while randomly changing the tones with the data slider.

In pattern play mode, there are 64 preset rhythm patterns (your basic bread-and-butter rhythms, nothing real exciting), and you can reprogram 32 of these yourself. A pattern signature can range from 1/4 to 8/4 (the denominator is constant), which gives you some rhythmic freedom, but don't expect too much, because a pattern is always one measure long.

The process of programming a D-20 rhythm pattern is much like working with a drum machine, although there is no visual display for each drum sound—something most of Roland's drum machines provide. You can quantize events as you enter them if you'd like; the available note resolutions are eighth-note, eighth-note triplet, sixteenth-note, sixteenth-note triplet, and thirty-second-note. If you don't like something in the pattern, you can erase the offending sound(s). We found it somewhat irritating that the stop button doesn't immediately end rhythm playback—the machine continues to the end of the current pattern. You use the D-20's sequencer to string the rhythm patterns together.

The Sequencer. First, the good points: In multitimbral mode, the D-20's eight-track sequencer can play up to eight polyphonic timbres and a rhythm pattern track. You can overdub note, program change, or volume and pan data to an existing track, then audition the results before merging the new data with the track. You can strip controller data from a track. After-the-fact quantization is available.

You can use track 8 to record rhythm events performed in real time, or you can use the rhythm track and string together rhythm patterns. Since rhythm patterns aren't restricted to 4/4 time, you can mix time signatures throughout the sequence, no matter what the original signature is. The sequencer can give you a two-bar lead-in or begin recording when you depress a key. Punch-in is automated.

If you have a more powerful sequencer for putting your music together, the D-20's sequencer can still be used in live performance, because it can record sequence data in real time from another sequencer on all eight tracks at once—a *real* time-saver.

Now, for the bad news: The D-20 will contain one, and *only* one, sequence at a time. Therefore, when you use the D-20 in live performance, you're forced to reload each new song, instead of having several in memory ready to go. You could separate each song with a blank measure and stop the sequencer between songs, but can you imagine the kind of trouble you'll have to go through if you need to compose a set of songs on the D-20's sequencer?

You must save a sequence to disk before you turn the instrument off, or you'll lose the sequence. Yes, this is how any software sequencer works, but most of the other synth/sequencer combos keep sequences in memory.

There are no copying functions, so you can't copy data from one track to another. Worse yet, you can't copy a group of measures to another part of the sequence. We wouldn't be so upset if there were individual sequences that we could string together in their preferred order. Another gripe: There is no single-event editing.

Unlike other synth/sequencer combos (such as the ESQ-1), a sequencer track on the D-20 can only play either one of its internal timbres or an external MIDI device on the proper MIDI channel—not both. Oh, and get this—to send a track over MIDI, you mute it. So you have an external synth hooked up to the D-20 while you're working on a sequence. You mute a track that's playing a D-20 timbre so you can hear another part, and what happens? If the channel matches up, you'll hear the track you muted on the other synth. Just as annoying, some important data is not saved to disk at all. Let's say you create a sequence that plays tracks 5 through 8 over MIDI channels 11 through 14. You finish the sequence and save it to disk. For your next sequence, tracks 5 through 8 play MIDI channels 1 through 4. When you reload the first sequence, tracks 5 through 8 won't sound right, because the D-20 doesn't store MIDI channels for each sequence. You'll need to change them manually if they're not the same for every song.

The manual doesn't tell you how to truncate a sequence if you let the sequencer run in record mode longer than you intended. Here's how to do it: While in record mode, erase the data in any track, starting at the measure *after* the last measure you'd like to hear play. When you return to play mode, the sequence will only run as far as the erase point. But what if you wanted to remove part of a single track without truncating the entire sequence? You can, but there's a second step involved. All the other tracks will still be there, but the sequencer won't play them. To fix this problem, stay in (or return to) record mode, start the sequencer, and let it run all the way to the sequence's end.

We have three more gripes, but they're minor: (1) Other than the audition capability of the overdub function, you can't check the results of quantization or punching in—the previous data is lost. (2) The sequencer won't stop playback until it reaches the end of the current measure. (3) MIDI Song Position Pointer is not implemented.

Data Storage. Anyone who has hassled with cassette tape data storage probably drools at the thought of an on-board disk drive. Random access can be a marvelous thing. Unfortunately, the D-20's disk drive won't work as a generic sys-ex data storage system.

While the D-20's drive does a good job of storing sequencer, rhythm, and sound data to 3½" disk, it's too bad you can't save individual sounds, rhythm patterns, or sequencer tracks. It takes almost 2½ minutes to format a disk, 40 to 103 seconds to save a sequence, and about 10 to 35 seconds to load a sequence. Saving all memory data, including a 500-bar sequence, takes about 125 seconds, while reloading takes just over 50 seconds.

You can use an optional RAM card for storing all programmable rhythm patterns and sound data. Although you can't store individual sounds to the card, you can simultaneously access sounds on a card and in internal memory.

Data dump commands are provided so that you can transfer the entire contents of memory from one D-20 to another, or to a sys-ex librarian.

Front Panel Interface & PG-10 Programmer. The D-20's front panel is laid out much more ergonomically than the D-50's, though the LCD window is smaller. Veteran LA programmers will be glad to find the D-20's increment and decrement buttons very near the data slider.

If you prefer to have access to more sliders, you should check out the PG-10. Its eight sliders and multiple pushbuttons can simplify programming. We're happy to report that, although it uses the D-20's MIDI in, the PG-10 provides a utility MIDI in for D-20 parameter data. This leaves the PG-10's primary MIDI in available for input from an external controller, in which case the PG-10's MIDI out acts as a thru. At the same time, the D-20 and PG-10 use each other's MIDI outs, so you'll need a thru box if you intend to use the D-20 as a master controller at the same time you're programming its sounds. Our only problem with the PG-10 was that its sliders have a relatively short throw, and there are no increment and decrement buttons. Imagine trying to zero in on one of up to 129 values using a slider that has only 1½ inches of movement. Increment and decrement buttons would solve the problem. Still, you may find that paying $399.50 for the programmer will make your life easier.

Conclusion. If you've been holding off on buying a D-50 because you were wishing it was multitimbral and had a built-in sequencer, disk drive, dynamic voice allocation, and drum sounds, the D-20 is definitely worth looking into. Our impression is that the sound quality isn't quite up to the D-50 standard, but for many people the extra features may tip the balance. We do wish we could be more enthusiastic about both the sequencer and the disk operating system. They're useable, but they're missing some important features that we'd expect to see on a quality product—features that the competition has had for a couple of years. We'd guess that Roland is hoping that the magic of LA synthesis will catapult the D-20 to the top of the do-it-all synth mountain, but they may be facing an uphill battle. ■

NEW EQUIPMENT DROUGHT?

KEYBOARD®

OCT. '88

ATTACK OF THE NEW AGE SYNTH HEROES

PLUS

FREE RECORD: T LAVITZ

LAST CHANCE: WIN A STEINWAY GRAND!

NEW COLUMN: MIDI CONTROLLERS

REVIEWS: ROLAND D-20, TURBOSYNTH, CASIO VZ-1

10

U.S. $2.95
Can. $3.95
U.K. £2.00

Display until 10/20

0 262936 0