SPECIAL ISSUE

KEYBOARD®

MAR. '89

SAMPLERS

THE COLD HARD FACTS

LAB AND LISTENING TESTS ON ALL THE TOP MODELS

MIDI-MALLET SUPERCHOPS

STRIDE PIANO CLINIC: PLUS SOLO TRANSCRIPTION

FREE RECORD INSIDE: READER COMPETITION WINNER

SAMP

THE REAL LOW-DOWN NITTY-GRITTY (AND WE DO MEAN GRITTY)

NEVER AGAIN. WE PROMISE. (At least, not until next time.) It seemed like such a simple idea, last summer when we first considered it: "Let's do a cover story with audio tests on all of the major samplers. In fact, let's make it a Special Issue. Yeah, that's the ticket! We can throw in a chart with sampler features, and maybe some listening tests too." We didn't *think* we were being especially naive. After all, we had done Special Issues before—on sequencers, on recording, on MIDI, on improvisation, on experimental music, on live performance, on basics, on the piano. . . . It's not that hard: You just put together a bunch of features on the same topic. Why should a Special Issue on samplers be any different?

We found out why. The first question we faced was, how exactly do you do audio tests on a sampler? Well, you, um, hook up some, you know, like, test gear or something, right? And you measure, like, noise and distortion and stuff. (Believe it or not, *Keyboard* staffers are not born with an innate, genetically inherited knowledge of the difference between a Hertz and a decibel.)

Realizing that we were in treacherous waters, we made half a dozen phone calls to folks in the industry—people like Tom Oberheim, Peter Gotcher and Evan Brooks at Digidesign, Dave Rossum at E-mu, and Marcus Ryle of Fast Forward Designs. They were more than generous with their time. But did the waters get less treacherous? No. We found that there are no accepted standards for how a sampler's audio performance ought to be measured. Every instrument designer devises his own tests, of course, in order to have some quality control over the product. But no two sets of tests are exactly alike. What's more, there are so many ways that a sampler's sound can fall short of perfection that by the time we did all possible tests, the instruments we were testing would be obsolete. So before we could design our own tests, we had to weigh the opinions of various experts on some highly technical subjects.

At length we made some tentative decisions about what the tests ought to measure. Now we were back on firm ground (or so we thought). But the task ahead was clearly massive: As many samplers as possible had to be sent to us by manufacturers, test equipment had to be gathered, and test tones had to be recorded into each sampler. *Then* we could start doing the tests.

Several weeks of work lay ahead. (An ominous chortle is heard in the back of the room.) Fortunately, former assistant editor Ted Greenwald, who moved to New York last year to pursue a career in session work, had a suitable hole in his schedule. By twisting his arm (actually, by offering him a pile of money, which works better over the phone), we convinced him to return to Cupertino to run the tests. We knew Ted would be the perfect choice, because, above all, he is meticulous about technical questions. If Ted was satisfied with the results of a test, we knew we could take the heat from an unhappy manufacturer.

Now that the ball was rolling, we announced the upcoming special issue to our advertisers. The following month, we announced it again. The month after that, we announced it a third time. We were starting to feel like the vaudeville MC who makes a grand gesture as the curtain rises, only to find that he is pointing to an empty stage.

No, the problem wasn't that Ted had the flu. Nor had the building burned down. The problem was that testing samplers turned out to be a whole lot more complicated than we expected. For details (lots of them), you can check out Ted's story, which begins on page 50. For now, we'll give just one brief example: On the very last instrument he tested, which happened to be the Ensoniq EPS, the numerical figure for noise was far higher than could be accounted for by the noise visible on the graphic display. Ted discovered that the EPS output contains a significant noise peak in the 50-100kHz range, which was off the right edge of the graph as he had originally set up the graph. Such noise isn't audible to human ears, so the test results would be more meaningful

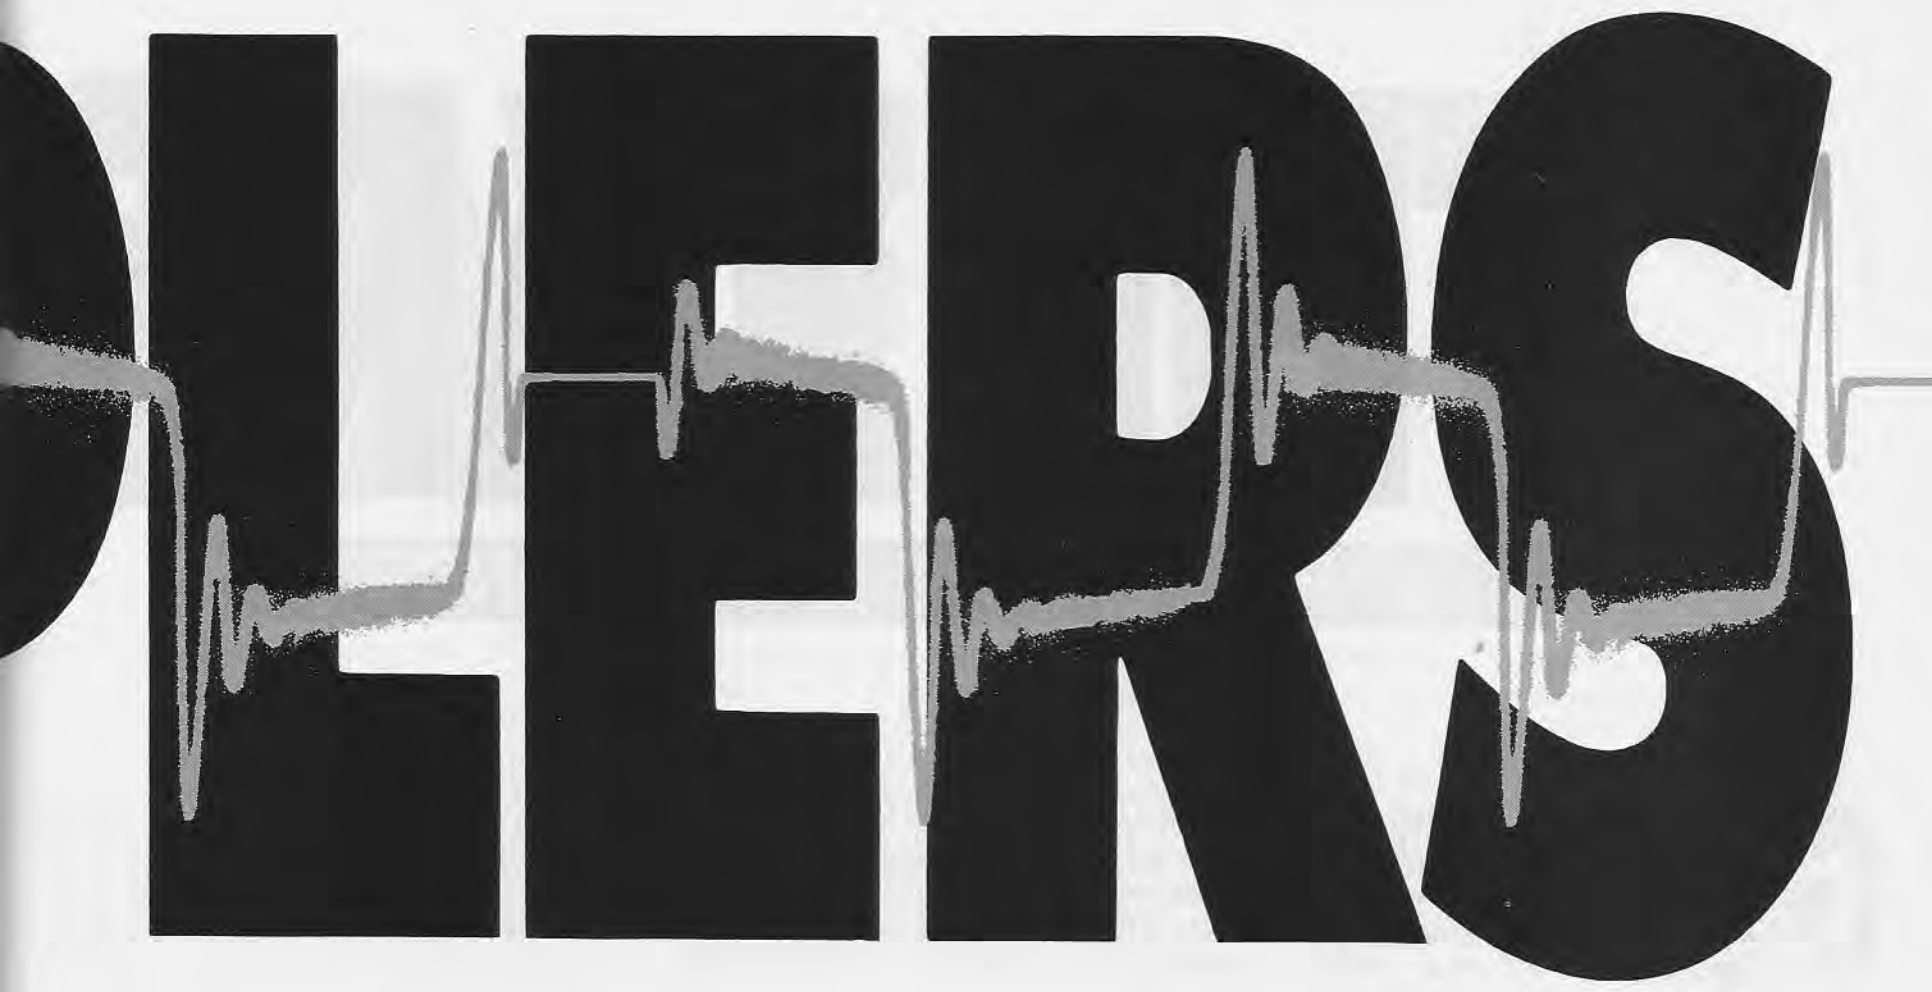

if they were band-limited to the 20kHz range. But in a fair test, you can't band-limit the output of one instrument; you have to limit them all. Which meant that Ted had to run *all* of his tests again from scratch.

Since we had upwards of a dozen samplers on hand, we figured it would be simple enough to load some identical sounds into all of them and do some blindfold listening tests. (The ominous chortle mutates into a guffaw.) Paul Wiffen, a freelance music journalist, sampling whiz, and sometime *Keyboard* columnist, spent a week in Cupertino preparing disks with samples for all the machines—no mean feat, considering that the operating systems of most samplers are among the most arcane, labyrinthine, and downright user-hostile pieces of software in the Land of Music.

Unfortunately, the results of our first battery of listening tests were inconclusive. We had problems matching the volume levels of the instruments, and we couldn't book enough time in the studio environment we had chosen. We needed to do a second set of tests in a room where we could set things up properly ahead of time. Paul, however, had to return to England. So one of our new assistant editors, Michael Marans, took over the preparations. It was *his* turn to learn all those operating systems. A couple of the owners' manuals were missing by this time, and a couple of the instruments had developed a tendency to crash every 15 minutes. Michael worked through the Christmas weekend getting everything shipshape. You can read about the final outcome of the listening tests on page 24.

Meanwhile, back at the ranch, assistant editor Mark Vail had been preparing The Chart. (Two-second snippet of shark music from *Jaws*.) "All you have to do," we explained to Mark, "is call all the manufacturers and ask them a bunch of questions, and write down their answers. Simple, right? Heh-heh-heh." Mark turned a little pale, but he squared his shoulders and manfully picked up the phone.

Two months later, we were still uncovering subtle ambiguities in the rough version of the chart. We found that various manufacturers have their own ways of describing how their machines work. Some questions are simple: Does it have a sequencer? Yes or no. Other questions are not nearly so simple: How are voices assigned to output jacks? Does the instrument actually have an analog lowpass filter, or does it use a digital filtering algorithm that only *mimics* an analog filter? And which technician at that company knows the answer to the question, and when will he be back from vacation?

We're not telling you all this so you'll feel sorry for us. (You can if you want to, though. We won't try to stop you.) We're doing it so you'll understand two things about this Special Issue: (1) There is a *lot* of information in it about samplers, information you will read nowhere else because no other magazine is foolish enough to put themselves through such contortions. And (2) on the afternoon when the production department came into the editorial offices with baseball bats and tore the copy out of our ink-stained, twitching fingers, there were still a few facts that we hadn't triple-checked, and maybe even a number somewhere with a suspicious decimal point.

If any errors did slip through, we're sure eagle-eyed readers will tell us about them, so keep your eyes on our Letters for the next month or two. Because one thing we do know: Our readers *care* about these things. Many of them make their living, or a portion of it, using samplers, so the differences between one model and another are not merely academic. Whether you're in the business of building and selling samplers, thinking about buying your first sampler, thinking about moving up to a better one, or simply wondering what all the fuss is about, this issue will have something for you.

We learned a tremendous amount in the process of putting it together, and we'd like to thank the many friends throughout the industry who made it possible. Maybe next time we'll try something just a little easier. Maybe a Special Issue on patch cords. —*JA* ■

THE EARS

SURPRISING RESULTS OF BLIND LISTENING TESTS

BY MICHAEL MARANS

GRAPHS. CHARTS. PLOTS. RESPONSE curves. Great stuff for the brain. But what about the ears? Can the ears distinguish an instrument with a noise floor of –96dB from one whose noise floor is –89dB? And can you really hear the difference between a sampler that reproduces a 10kHz sine wave with .5% distortion and one that distorts the wave by .7%? Sometimes the technical subtleties of instruments have little to do with their real-world musical applications, so in order to present a more complete picture we decided to include some subjective evaluations of the instruments we were testing. We set out to conduct a little experiment—make that a *huge* experiment—in the form of critical listening tests.

The premise was simple enough; sample the same eight or ten sounds into all of the instruments and have a panel of experts (who just happened to be *Keyboard* staffers and a couple of local musicians) judge them for fidelity, noise, musical usefulness, and so on. Ace sampler Paul Wiffen was flown in to tackle the job of sampling the sounds and conducting the tests. As usual at *Keyboard*, Paul was given less time to accomplish his task than needed by any mere mortal, but met the deadline through a combination of superhuman effort and no sleep. L.A. session engineer Joey Wolpert was brought in to set up the listening environment and ensure that each instrument would be accurately monitored. He brought with him a mixing console intended specifically for keyboard instruments, and designed to

TEST #1	UNPROCESSED/ ORIGINAL PITCH				
	OVERALL	THIN/FAT	NOISEY/ CLEAN	DULL/ BRIGHT	PROCESSED/ TRANSPARENT
CASIO FZ-10M	8.56	8.37	9.00	8.37	8.50
E-MU EMULATOR THREE	7.96	7.75	8.25	8.00	7.87
AKAI S1000	7.56	7.87	7.12	7.37	7.87
SIMMONS SDX	7.25	7.50	8.75	5.87	6.87
SEQUENTIAL PROPHET 2002	5.90	6.50	5.50	6.75	4.87
ROLAND S-550	5.81	6.50	4.50	7.00	5.25
ENSONIQ EPS	5.71	6.12	5.87	6.12	4.75
KORG DSM-1	5.46	6.00	6.50	4.62	4.75
ROLAND S-330	5.37	6.00	4.75	5.87	4.87
AKAI S950	4.62	5.87	3.25	5.75	3.62
YAMAHA TX16W	4.50	5.75	3.75	4.75	3.75
AKAI S900	4.31	5.50	2.12	5.00	4.62
E-MU EMAX	4.18	5.62	3.12	4.37	3.62
KURZWEIL 250	3.25	5.00	1.50	4.12	2.37
ENSONIQ MIRAGE MULTI-SAMPLER	1.53	2.50	1.00	1.62	1.00

TEST #2	TRANSPOSITION
AKAI S1000	7.42
E-MU EMULATOR THREE	7.28
AKAI S950	6.71
YAMAHA TX16W	6.57
ROLAND S-550	6.14
E-MU EMAX	5.57
SEQUENTIAL PROPHET 2002	5.28
SIMMONS SDX	5.14
KURZWEIL 250	5.14
ROLAND S-330	5.14
AKAI S900	4.28
KORG DSM-1	3.14
ENSONIQ EPS	2.85
CASIO FZ-10M	1.71
ENSONIQ MIRAGE MULTI-SAMPLER	1.14

TEST #3	CHORDS
AKAI S950	7.00
AKAI S1000	6.42
YAMAHA TX16W	6.28
E-MU EMULATOR THREE	6.00
ROLAND S-550	5.71
SEQUENTIAL PROPHET 2002	5.57
SIMMONS SDX	5.42
ROLAND S-330	5.28
E-MU EMAX	5.00
KURZWEIL 250	4.85
AKAI S900	4.71
KORG DSM-1	4.28
ENSONIQ EPS	3.28
CASIO FZ-10M	2.71
ENSONIQ MIRAGE MULTI-SAMPLER	2.00

HAVE IT

A SPECIAL REPORT

be as transparent and noiseless as a console can be.

Despite our best efforts and intentions, it soon became apparent that there were inherent problems with the test procedures. For example, although the instruments were set to play back at identical volume levels, the individual samples were all slightly different in volume, and this meant readjusting the mixer for each sample. (Volume level turned out to be a vital factor when comparing the fidelity of the instruments.) When MIDI reared its ugly head and a little human error popped in, we concluded that it was imperative to conduct. . .

The listening tests, round two. We had discovered in the first test that strings and brass were not the most telling of sounds. Virtually all of the instruments were adept at reproducing these sounds, but drums and piano did a better job of exposing each sampler's response characteristics. As a result, these were the sounds we chose for the second test series. Additionally, the number of tests was reduced, and a different person was assigned the task of sampling the sounds into the instruments—me. (While *Keyboard* readers were out celebrating the holidays, I was staring into a stack of LCD screens and computer monitors.) There were two major changes in the test procedures: First, we would sequence all playback (to eliminate human error), and second, we would listen both to raw samples and to samples that used analog processing to clean up noise. This would allow us to judge instruments not only on their abilty to accurately reproduce a given sound source, but also on their ability to make a "bad" raw sample musically useful. Although the tests were not completely devoid of minor hitches, we were able to arrive at results that we feel are eye-opening and informative. When combined with the technical tests and the features chart, they provide a great deal of insight into the capabilities of each sampler.

ABOUT THE TESTS

The first test was designed to rate an instrument's ability to accurately reproduce a sound source. All samples were

TEST #4	PROCESSED				
	OVERALL	THIN/FAT	NOISEY/ CLEAN	DULL/ BRIGHT	PROCESSED/ TRANSPARENT
AKAI S1000	8.10	7.42	8.57	8.42	8.00
E-MU EMULATOR THREE	7.64	7.28	8.28	7.42	7.57
ROLAND S-550	6.89	7.14	8.00	7.00	5.42
KURZWEIL 250	6.78	6.42	7.14	7.14	6.42
AKAI S900	6.75	6.00	8.28	6.28	6.42
KORG DSM-1	6.5	7.00	7.71	5.85	5.42
AKAI S950	5.96	5.28	7.57	5.85	5.14
SIMMONS SDX	5.75	5.42	7.00	6.00	4.57
ROLAND S-330	5.74	6.28	6.00	6.28	4.42
SEQUENTIAL PROPHET 2002	5.64	6.14	6.28	5.42	4.71
E-MU EMAX	5.60	5.57	6.00	5.42	5.42
YAMAHA TX16W	5.42	5.71	5.42	5.42	5.14
CASIO FZ-10M	4.46	5.14	4.00	4.57	4.14
ENSONIQ EPS	4.10	4.85	3.57	4.42	3.57
ENSONIQ MIRAGE MULTI-SAMPLER	2.32	5.28	1.14	1.71	1.14

	CLEAN(COMBINED PROC./UNPROC.)
E-MU EMULATOR THREE	8.26
SIMMONS SDX	7.87
AKAI S1000	7.84
KORG DSM-1	7.10
CASIO FZ-10M	6.78
ROLAND S-550	6.25
SEQUENTIAL PROPHET 2002	5.89
AKAI S950	5.41
ROLAND S-330	5.37
AKAI S900	5.20
ENSONIQ EPS	4.72
YAMAHA TX16W	4.58
E-MU EMAX	4.56
KURZWEIL 250	4.32
ENSONIQ MIRAGE MULTI-SAMPLER	1.07

	OVERALL RATINGS
AKAI S1000	7.34
E-MU EMULATOR THREE	7.22
ROLAND S-550	6.14
AKAI S950	6.07
SIMMONS SDX	5.89
YAMAHA TX16W	5.69
SEQUENTIAL PROPHET 2002	5.60
ROLAND S-330	5.38
E-MU EMAX	5.09
AKAI S900	5.01
KURZWEIL 250	5.00
KORG DSM-1	4.84
CASIO FZ-10M	4.36
ENSONIQ EPS	3.99
ENSONIQ MIRAGE MULTI-SAMPLER	1.74

taken at an instrument's highest sampling rate. A sequence was created that played each sample one at a time, and then as part of a rhythm pattern. Samples were played back at original pitch at velocity 127, and no processing was used. The judges (who did not know which instrument they were listening to) were asked to rate the sounds in four categories: thin versus fat, noisy versus clean, dull versus bright, and processed versus transparent. Each category used a rating scale of one to ten, with one being thin, noisy, dull, or processed sounding, and ten being fat, clean, bright, or free from any coloration. The following drum sounds were used for this test: kick,

USER-HOSTILITY A GUIDE TO OPERATING SYSTEMS

By Michael Marans

WHEN YOU GET A NEW SAMPLER out of the box, the first thing you're likely to want to do, after playing the factory disks for an hour or so, is record a sound, assign it to the keyboard and/or a MIDI channel, maybe add an amplitude envelope, and play it. Sounds simple enough when you describe it in those terms, but in the course of preparing our listening tests, we found it anything but simple. With a few notable exceptions, most samplers turn these basic functions into labyrinthine nightmares.

Of course, most operating systems *do* make sense once you learn them. The fact that it might take twenty button pushes to do something that another sampler would allow you to do with four button pushes may be annoying, but it isn't terribly significant once you've got the button pushes memorized. And there are sometimes (sometimes) advantages to be gained from a multi-level operating system. But first impressions reveal a lot. While we didn't have time to master every nuance of every instrument, we did spend enough time with them to get thoroughly frustrated. Our experience left little doubt as to why there are so many third-party sound developers. In the interest of promoting the mental health of our readership, we present a quick look at best and the worst, and detail a few of the more unfortunate systems we encountered.

Without question, the most confusing, most aggravating, most impossible to decipher, most, most *aarrgh* system can be found on the Yamaha TX16W. By the way, we're referring to their new, improved, easy-to-use version 2.1. Due to its multitude of edit buffers, internal memories, performance setups, voices, waves, and layers, the TX16W is perfect for playing disks created by *someone else*. But even that is a less than desirable prospect, because the unit sports the longest disk load time of any of the instruments we tested, and requires multiple disks for booting, editing, and playback.

Second runner-up to the Yamaha is the Kurzweil 250 (which also wins the Laurens Hammond Memorial B-3 Portability Award). Actually, sampling was a breeze, albeit full of numerous button pushes, but we couldn't make heads nor tails out of the instrument/keyboard setup architecture. We were finally able to get it working after a call to the engineer who wrote the system. New software (version 6.0) has been written that we're told simplifies the procedures in this area. Of course, you'll need a Macintosh computer if you want to save and load your samples. If all you want to do is play back the factory presets, you'll never encounter the difficulties that we had to wrestle with.

The Korg DSM was less than a joy to work with. The voice architecture, in our

AVAILABLE AT YOUR LOCAL MUSIC STORE

Power in Your Pocket

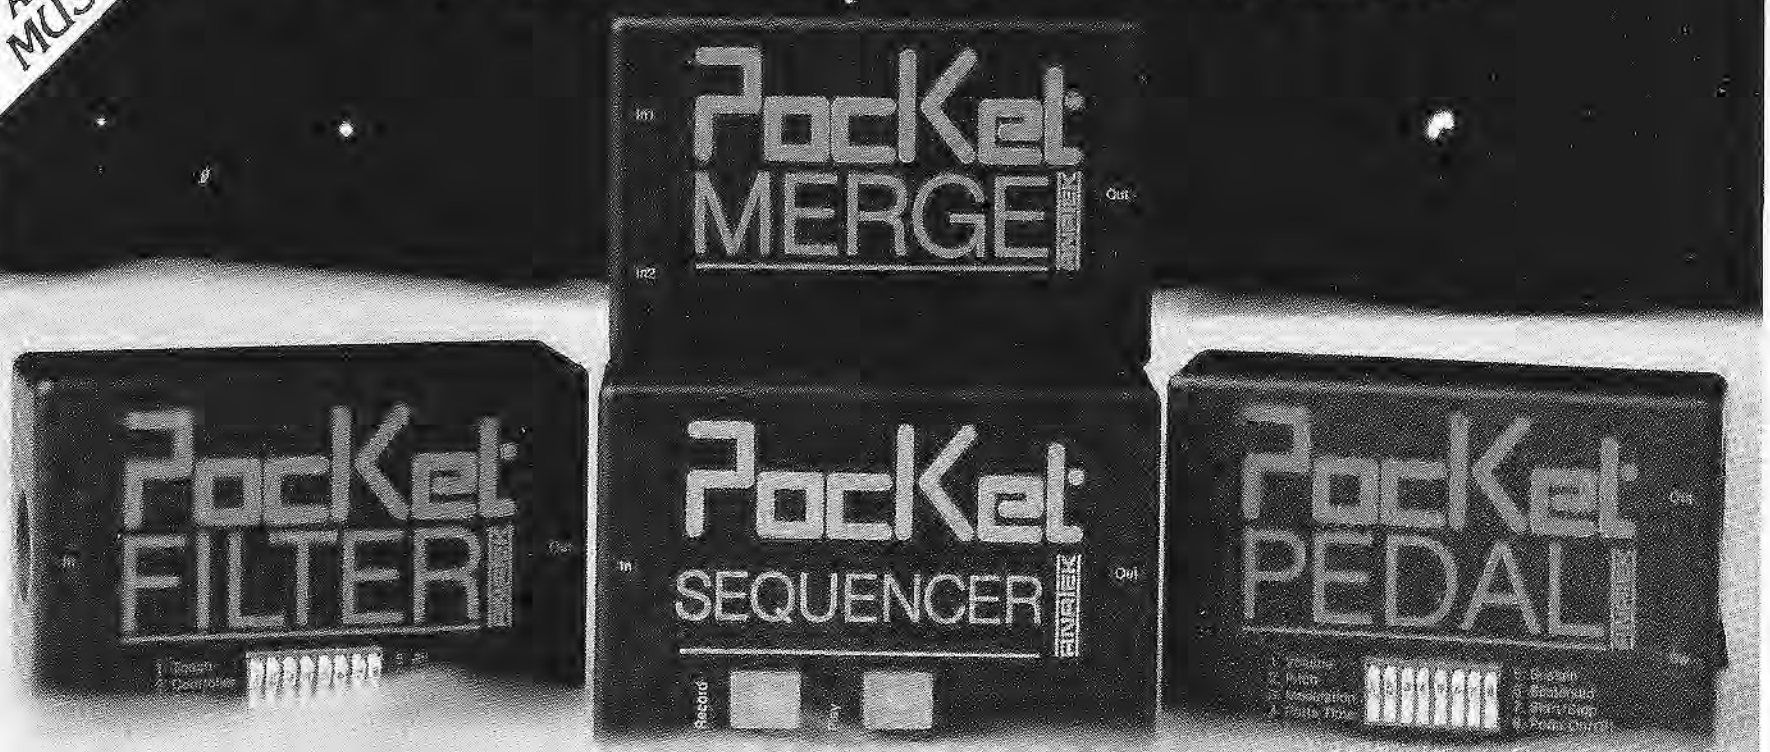

Pocket Products™ — affordable MIDI accessories that require no batteries or power supply. These user friendly accessories will be a necessary part of any MIDI system.

ANATEK
Music Products Division
400 Brooksbank Avenue
North Vancouver, B.C. V7J 1G9
Fax (604) 980-2722 Tel (604) 980-6850

© 1988 Anatek Microcircuits Inc.

snare (with reverb), closed hi-hat, open hi-hat, crash cymbal, and low tom. Each of the sounds offered a particular challenge to the instruments. The decays of the snare reverb and the tom would show noise, the cymbal would show frequency response, and the kick, snare, and hats would demonstrate the ability to reproduce sharp attack transients.

Test number two rated a sampler's ability to transpose a sound up a half-step (the limit of certain samplers at their highest sampling rate) and down a minor seventh. We chose a piano sample for this test since it is one of the least forgiving sounds when transposed. The samples were played at full velocity and no processing was used.

Continued on page 137

opinion, is needlessly complicated. By the time we had assembled our sound (wave, multi-sample, program, timbre, combination), we were too tired to play. Bunches of button pushes, limited maximum sample length, and the necessity to create both a work disk and a performance disk put this one way down on our list.

The Roland S-550 was fairly straightforward to learn, though we definitely recommend using a video monitor. This instrument also suffers from having too many levels. Tone, subtone, tone map, performance setup, who needs it? Trying to determine where to assign what (not to mention bouncing around among the numerous main menus, submenus, and command menus in the process), coupled with an excessive number of mouse clicks, left us with the desire to take a nice Hawaiian vacation and learn to play the ukulele. Interestingly enough, a few of the identical functions on the less expensive S-330 were easier to use. Both machines limit the maximum sample length, and have the annoying habit of adding a strange buzzing sound during the decay portion of a sample whenever the TVA (Time Variant Amplifier) envelope is used. (Yamaha's digital envelopes exhibited a similar phenomenon.)

The Simmons SDX is the only sampler we tested that had a built-in video monitor screen. Overall, we found the display to be clear and concise, and operation of the unit was simple and easy. However, due to the heirarchical nature of the software, it was often necessary to plow through several menus to access a particular function, and wind back out to get to another. The default parameter values, while acceptable for drum sounds, have no direct correlation to pitched sounds, and only limited global parameters are available.

The Casio FZ-10M suffers from the same heirarchical menu problems. While it is easy to use and quite logical, it wins the award for most button pushes required. The user must also operate it with two hands, as the cursor, enter, and escape buttons are on one side of the machine, while the increment/decrement buttons and the data slider are on the other. The waveform display was very useful (especially since we never did get the knack of reading the level meter), and allows a lot of information to be simultaneously displayed. The display flicker, however, was enough to force a break in the action while we ran out for some Murine.

Speaking of incessant button-pushing, before you buy an Ensoniq EPS, make sure that getting calluses on your fingers won't interfere with your playing. We also noticed that many of the global settings are not retained when you power down, and must be re-entered each time you play. As for their other entry in our tests, the Mirage, what can you say about an instrument that requires all data to be entered in hexadecimal code?

The five easiest instruments to learn came from Akai (S900, S950, S1000) and E-mu (Emax, EIII). The operating systems designed by these two manufacturers, though conceptually quite different, are logical, clear, concise, and easy to use. If you're searching for an instrument that you can start programming quickly, with barely a glance at the manual, these are the ones to look at first. ■

smpte gets smarter.

J.L. Cooper's PPS-100 is a SMPTE/EBU to MIDI Synchronizer and Event Generator with smarts. The PPS-100 converts all SMPTE formats to MIDI Sync and Song Position Pointer with programmable tempo changes. The PPS-100 also generates MTC (MIDI Time Code), DIN sync, Direct Time Lock, and PPQN sync.

Need to generate time code as well? The PPS-100 generates all formats, starting at any time. And the PPS-100's "event generator" is even smarter! Any kind of MIDI commands (notes, program changes, even systems exclusive messages) can be output at any SMPTE time.

Two independent on-board relays and two pulse outputs may be programmed to turn on and off at any SMPTE time. Unlimited uses include auto punch-in and automated effects bypass.

Optional software "desk accessory" for the Atari ST and Macintosh allows for cue list entry of events.

Smart price, too. Just $595. Optional disk only $49!

Ask for it now at your music and pro audio dealers!

JL COOPER ELECTRONICS 13478 Beach Ave., Marina Del Rey, CA 90292 (213) 306-4131

Continued from page 27

Judgments were made on distortion (including aliasing) and tonal coloration. Ratings were made on a scale of one to ten, ten being absolutely no distortion or coloration.

The third test also used the piano sample, unprocessed and at full velocity, to determine an instrument's ability to play multiple, simultaneous pitches. This (in theory) should enable us to distinguish the single DAC instruments from those with multiple DACs. (It's like having 16 guitar players plug into the same amplifier versus each one having their own system.) The judges listened for attack smearing, phase cancellation and/or distortion, and tonal coloration. You guessed it; the same one-to-ten ratings scale was used.

The fourth and final test used the piano sample, but this time it had been processed with analog parameters. Primarily, an amplitude envelope was used to eliminate noise on the sample's decay, but in certain instances a small amount of filter envelope was also used. The object was to use the minimum amount of processing required to clean up a sample, but at the same time to use as much as necessary to eliminate all noise. Keyboard velocity was assigned to sample playback level, and on each instrument a curve was selected that provided the most even response. A combination of the sequences from tests two and three was used, but assorted velocity values were added to evaluate tone and envelope integrity at multiple volume levels. The same four ratings categories as in test number one were used, and in case you forgot, one is worst, ten is best.

ABOUT THE CHARTS

Unprocessed. This chart lists the instruments in order of their rating—highest to lowest—based on each sampler's performance in the raw sample fidelity test. This rating, shown in the first column, was derived by averaging the scores in columns two through five. In our averaging, we gave equal weight to each of the four categories; you may find one category much more important than another depending upon your specific applications. For example, if you are using your sampler strictly to play drum sounds in live performance situations, you may be much more concerned about the ratings for fat and bright than those for noise. Of course, placing different weights on different categories could ultimately affect an instrument's position in the ratings. It is also important to realize that this particular test used drum samples, and the results are meaningful only in the context of an instrument's ability to cleanly reproduce the specific sound sources we chose for the test. For example, the winner of this test, the Casio FZ10M, scored poorly in test number four when it was required to reproduce a piano sound. It is important to examine each category in each test in order to gain a realistic appreciation of a particular sampler's strengths and weaknesses.

Processed. This chart is should be interpreted in the same manner as the unprocessed chart. This test used a piano sample, which some of the samplers (K250, S900) had difficulty reproducing in its raw state without adding noise. Processing the samples with a small amount of amplitude envelope greatly improved their performance. Some of the samples also acquired noise artifacts, as explained in the next section, which undoubtedly affected their relative standings.

Transposition & Chords. The transposition test provided two challenges to the instruments; one, to reproduce a piano sample, and two, to transpose it over a wide range. It is usually necessary to pitch-shift a sample across a portion of the keyboard, and in our less-than-perfect world this often introduces distortion. Five of the samplers exhibited peculiar problems. The Casio FZ-10M and the Roland S-330 each introduced clicks into the sound. (This could have been caused by aliasing of supersonic frequencies which may have escaped the output filter of our source CD player and the input filters of the instruments. In each case, we tried taking the sample several times and always got the same clicks. It is likely, however, that different sources will produce different results.) The SDX added no clicks, but developed distortion at approximately the same point in the sample as the Casio. The DSM sample also contained distortion, which repeated rhythmically as though it had been looped. Lastly, the EPS twisted the sample's initial attack quite markedly.

Clean. These figures were derived by averaging each instrument's noisy/clean score in both the unprocessed and processed tests.

Overall Ratings. This chart lists from best to worst the average score received by each instrument in all four tests (not including the separate "Clean" results).

The test results contained many surprises. Instruments that were expected to do poorly occasionally ended up winners. Conversely, expected winners often found resting places in the lower depths of the ratings. Judging from the "oohs," "ahhs," and "you're kiddings" heard when the identities of the instruments were revealed to the judges (the order of playback was shuffled for each test), it's safe to say that some of the pre-conceived notions you may have regarding the sonic qualities of a particular sampler are likely to be less than accurate. If you are thinking about adding a sampler to your keyboard setup, you owe it to yourself to listen—and carefully—to a number of different sounds in a qualified, reference listening environment. Don't let a store salesperson play sounds for you through a mediocre system, and by all means, request that you be allowed to hear the instrument (without reverb) in a room that doesn't have a bunch of heavy metal guitarists trying out the latest stack of Marshall amplifiers.

You may find that a number of samplers sound equally good—we did. (A lot of the instruments in this chart are separated by tenths of points.) Naturally, you'll want to take a number of things into consideration when purchasing a sampler: features, library, manufacturer support, operating system user-friendliness, and of course, price. Only by combining all of these factors can you get a clear picture of the sampler that's right for your needs. ■

THE TANGLED THREAD

SAMPLER FEATURES & TERMINOLOGY EXPLAINED

BY JIM AIKIN

A CAMERA DOES ONLY ONE thing: It captures images on film. But around that seemingly simple idea a whole constellation of interlocking elements has gathered—film type, lens type, shutter speed, depth of focus, and so on. You can take snapshots without knowing anything about any of that stuff. Just point and click. But don't expect to get professional-quality results.

An even more bewildering profusion of options and terminology has grown up around the audio equivalent of the camera. And the same *caveat* applies. You can buy a sampler and play it without knowing anything about bit resolution, looping, or velocity cross-switching, but you'll never get the most out of your machine (unless you can afford to hire a technician to do your sampling). And if you should want to get better musical results, you won't have any way of determining whether you can get them from your current sampler, or whether you need to upgrade. If you do need to upgrade, which features are important to you, and which can you make compromises on?

In this article we'll take a look at the major features found on today's samplers. In the process, we'll explain the items listed in the chart that begins on page 30. Some of the concepts are vital aspects of sampler design, and some of them are much less important. Some of the choices that a manufacturer makes when designing a sampler have hidden tradeoffs, and we'll try to shed some light on these as well. Audio and listening tests may reveal the most important aspects of sampler performance, but a great-sounding sampler won't be much use to you if it doesn't have the special features you need. The chart allows you to make direct comparisons between samplers, which should make your purchasing decisions easier. For an explanation of the abbreviations used in the chart, see the key at the bottom of each chart page.

POINT-BY-POINT COMPARISON CHART PAGE 30

Most of the facts and figures in the chart were provided by the manufacturers. While we double-checked odd-looking items and got independent verification where possible, it's always possible that two companies might use the same term in different ways, or that a company representative misunderstood one of the questions we asked. While we made every effort to insure that the chart is accurate

	CURRENT OPERATING SYSTEM	LIST PRICE	ON-BOARD MEMORY (WORDS/ BYTES)	MEMORY EXPANSION	BIT RESOLUTION (A/D, INT., D/A)	No. OF VOICES	No. OF SOUNDS ACTIVE AT A TIME	No. OF SAMPLES IN MEMORY
AKAI S612*	ROM	$795	43K words/64K bytes	no	12/12/12	6	1	1
AKAI S700\X7000	SMX-007: 2.1.ASK-70: 1.6.	S700: $1,399.95 X7000: $1,449.95	512K words/750K bytes	Expandable to 750K words (1 meg bytes)	12/12/12	6	16	16
AKAI S900*	1.2	$2,499.95	512K words/750K bytes	no	12/12/12	8	32	32
AKAI S900 w/MARION SYSTEMS MS-9C BOARD	any S900 operating system	$995, plus Akai S900	same as S900	same as S900	12-bit mode: 12/12/12 16-bit mode: 16/12/16	same as S900	same as S900	same as S900
AKAI S950	1.0	$2,499.95	512K words/750K bytes	Expandable to 1.5 meg words (2.25 meg bytes) in 512K-word steps	12/12/12	8	99	99
AKAI S1000	1.1	$5,999.95 to $12,050 (w/8 meg RAM & 40-meg HD)	1 meg words/2 meg bytes	Expandable to 4 meg words (8 meg bytes) in 2-meg-word steps	16/16/16	16	200	200
AKAI MPC60	2.0	$3,299.95	512K words/750K bytes	Expandable to 1 meg words (1.5 meg bytes)	16/12/16 (with proprietary encoding)	16	34	34
CASIO FZ-1	ROM-B	$2,495	512K words/1 meg bytes	Expandable to 1 meg words (2 meg bytes) with 512K-word exp. board	16/16/16	8	64	64
CASIO FZ-10M	ROM-B	$2,495	1 meg words/2 meg bytes	no	16/16/16	8	64	64
E-MU EMAX	3.2/1.1HD. SE: 1.1 SE, 1.1 SE HD	$2,995 to $4,295. Rack: $2,795 to $4,095.	512K words/512K bytes	no	12/8/12 (w/internal data compression)	8, 16 in dual mode	176	406
E-MU SP-1200	1.2	$2,995	256K words/384K bytes	no	12/12/12	8	32	32
E-MU EMULATOR II*	3.1	$7,995	512K words/512K bytes	Expandable to 1 meg words (1 meg bytes)	8/8/8 (w/data compression)	8	122	100
E-MU EMULATOR THREE	2.0	$9,995 to $13,995	4 meg words/8 meg bytes	Expandable to 8 meg words (16 meg bytes)	16/16/16	16	1,496	100
ENSONIQ EPS	2.3	$1,995	256K words/480K bytes	Expandable to 512K words (896K bytes)	13/16/16	12@52kHz 16@39kHz 20@31.2kHz	127	1,016
ENSONIQ MIRAGE DSK/ MULTI-SAMPLER*	3.2	DSK: $1,295 Multi-Sampler: $1,195	128K words/128K bytes	no	8/8/16	8 (dual oscillators)	16	16
FAIRLIGHT CMI SERIES III	7.30	$65,000 to $120,000	7 meg words/14 meg bytes	Expandable to 14 meg words (28 meg bytes) in 1-meg-word steps	16/16/16	16	1,024	1,024
KORG DSS-1*	36. Mem exp: 29.	approximately $2,400	256K words/384K bytes	Expandable to 2 meg words (3 meg bytes) in 256K-word steps	12/12/12	8 (dual oscillators)	31	256
KORG DSM-1	14	approximately $2,600	1 meg words/1.5 meg bytes	no	12/12/12	16	64	512
KURZWEIL 250 w/SAMPLING OPTION	V6	$9,995 to $15,000	512K words/640K bytes	Expandable to 2 meg words (2.5 meg bytes)	10/10/10	12	528	256
NEW ENGLAND DIGITAL SYNCLAVIER 3200/9600	1.0	$50,000 to $250,000+	2 meg words/4 meg bytes	Expandable to 48 meg words (96 meg bytes) in 2- or 8-meg-word steps	16/16/16	4 to 96	352	limited only by RAM
OBERHEIM-ECC DPX-1	2.2	$2,095	667K words/1 meg bytes	no	NA/12/12	8	same as emulated sampler	same as emulated sampler
ROLAND S-10*/S-220	ROM 2.07	S-10: $1,295 S-220: $1,395	128K words/192K bytes	no	16/12/16	S-10: 8 S-220: 16	S-10: 4 S-220: 8	4
ROLAND S-330	1.02	$2,295	500K words/750K bytes	no	16/12/16	16	32	32
ROLAND S-50	2.02	$3,495	500K words/750K bytes	no	16/12/16	16	32	32
ROLAND S-550	1.12. w/ HD: HD5IF, ver. 1.02	$3,495	1 meg words/1.5 meg bytes	no	16/12/16	16	64	64
SEQUENTIAL PROPHET 2000*/2002*	4.2	2000: $2,999 2002: $2,799 to $2,999	512K words/768K bytes	Expandable to 1 meg words (1.5 meg bytes)	12/12/12	8	16	16
SEQUENTIAL STUDIO 440	2.10	$4,665	500K words/750K bytes	no	12/12/12	8	32	32
SEQUENTIAL PROPHET 3000*(YAMAHA P3000*)	2.04	$3,995	1 meg words/2 meg bytes	no	16/16/16	8	433	433
SIMMONS SDX	3.0	$8,000 to $14,000	1 meg words/2 meg bytes	Expandable to 4 meg words (8 meg bytes) in 1-meg-word steps	16/16/16	16	144	999
WAVEFRAME AUDIOFRAME	3.0	$44,400 to $140,000	1 meg words/2 meg bytes	Exp. to 60 meg words (120 meg bytes) in 3-, 7-, or 14-meg-word steps	16/16/18	16 to 256	1,024	1,024
YAMAHA TX16W	2.0	$2,895	1 meg words/1.5 meg bytes (includes 250K-word optional filter tables)	Expandable to 4 meg words (6 meg bytes) in 1.5-meg-word steps	12/12/12	16	256	512

Key: AES/EBU, Audio Engineering Society/European Broadcast Union; DAC, digital-to-analog converter; DVA, dynamic voice allocation; Fxd, fixed voice allocation; giga, billion; Hz, Hertz (cycles per second); HD, hard drive; IO, individual output; K, kilo (thousand); meg, million; mono, monophonic; poly, polyphonic; PVA, programmable voice allocation; RAM, random access memory; ROM, read-only memory; vce, voice; *, out of production.

No. OF OUTPUTS & TYPE (MONO, POLY)	No. OF DACs	PITCH SHIFT TECHNIQUE	VOICE ALLOCATION/ ASSIGNMENT MODES
1 poly, single timbre	1	variable playback rate	Output assignment Fxd
Single mix. Optional cable for 8 mono IOs.	1	variable playback rate	DVA; output assignment Fxd
Single mix, left & right, 8 mono IOs.	8	variable playback rate	DVA: each sound assignable to any output
same as S900	8	same as S900	same as S900
Single mix, left & right, 8 mono IOs.	8	variable playback rate	DVA: each sound assignable to any output
Left & right, 8 poly IOs, single mix send, stereo returns.	8	real-time mathematical decimation interpolation	DVA & PVA: each sound independently assignable to any one or more IOs
Left & right, 8 poly IOs, mono send, stereo returns.	1	real-time mathematical decimation interpolation	DVA:voice won't cut itself, except hi-hat; if all used, 1 w/least time left stolen
Single mix, 8 mono IOs.	2	variable playback rate	DVA (mix) & PVA (IOs)/ each sound assignable to any one or more IOs
Single mix, 8 mono IOs.	2	variable playback rate	DVA (mix) & PVA (IOs)/ each sound assignable to any one or more IOs
Single mix, left & right, 8 mono IOs.	8	drop/add sample with oversampling by 20 times	DVA & PVA: all sounds fixed to stereo outputs, assignable to IOs
Single mix, 8 mono IOs.	1	drop/add sample	DVA&PVA:each snd. assignable to 1 or 2 IOs; plug in IO removes voice from mix
Single mix, 8 mono IOs.	8	variable playback rate	DVA&PVA:each snd. assignable to 1 or 2 IOs; plug in IO removes voice from mix
Single mix, left & right, 16 mono IOs.	16	variable playback rate	DVA&PVA:each snd. assignable to 1 or 2 IOs; plug in IO removes voice from mix
Left & right. Optional 8 poly-IO module.	1	drop/add sample with linear interpolation	DVA: any voice or inst. assignable to IOs by kbd range, sample, or track
DSK: left & right. Multi-Sampler: single mix.	1	drop/add sample with zero-order interpolation	DVA; output assignment Fxd
Single mix, 16 mono IOs	16	variable playback rate	PVA: graphic assignment grid on CRT screen
2 single mixes	1	variable playback rate	PVA: layering or stacking possible; output assignment Fxd.
Single mix, 16 mono IOs.	1	variable playback rate	DVA or PVA: 5 modes, inc. no voice stealing, specific assignment, all IOs same
Left & right. Optional 12 mono IOs.	12	variable playback rate with tracking filter	DVA & PVA: any sound assignable to 1 or more IOs & stereo outputs
Left & right, 16 poly IOs.	3200: 32 9600: 192	variable playback rate	DVA & PVA: any sample assignable to 1 or more IOs
Single mix. Optional 8 IOs; mono or poly (same as emulated sampler).	8	variable playback rate	same as emulated sampler; sounds assigned as on disk
S-10: single mix. S-220: 4 poly IOs.	1	drop/add sample with differential interpolation	S-10: Fxd. S-220: PVA; 5 assignment modes define polyphony
Single mix, 8 poly IOs.	1	drop/add sample with differential interpolation	DVA & PVA: 22 modes; all sounds always at mix, assignable to 1 or more IOs
Single mix, 4 poly IOs.	1	drop/add sample with differential interpolation	Fxd: 8 modes; all sounds always at mix, assignable to 1 or more IOs
Single mix, 8 poly IOs.	1	drop/add sample with differential interpolation	DVA & PVA: 22 modes; all sounds always at mix, assignable to 1 or more IOs
Left & right. Optional 8 mono IOs.	8	variable playback rate	DVA or Fxd: each sound fixed to stereo outs; assn. to 1 or more optional IOs
Single mix, left & right, 8 mono IOs.	8	variable playback rate	DVA out stereo outputs; PVA to IOs
Single mix, left & right, 8 mono IOs.	8	variable playback rate	DVA or PVA: any sample assignable to any IO
Stereo, 16 mono IOs.	16	variable playback rate	DVA & PVA: each sound assignable to 1 or more IOs
Analog: 8 poly. Options to 64 poly. Dig.: 4 stereo for AES/EBU; 4 mono,all stnds	8 to 64 (1 per analog output)	interpolation	DVA & PVA: any sound assn. to 1 or more IOs; all sounds assignable to 1IO
Left & right, 8 mono IOs.	2	drop/add sample with linear interpolation	PVA: sounds assignable to either or both stereo outputs, or to IOs

Continued on page 32

SAMPLER FEATURES

and uniform, we can't guarantee that every instrument you'll encounter in a music store will have exactly the features indicated by the chart. Instruments with older operating systems sometimes sit on the shelves for months after an upgrade is available. Which brings us to . . .

Current Operating System. The operating systems of most samplers are upgraded at some point during the life of the instrument. The figure in this column tells which version is being described in the rest of the chart.

Some operating system upgrades are in the form of ROMs, while others come on disk. To install a new ROM, the instrument's case must be opened up, an operation that is best left to a qualified service technician. Disk operating systems can be loaded by the user, but must be loaded every time the instrument is turned on, which takes longer. A few samplers (notably the Ensoniq Mirage) have third-party operating systems available that will do some special tricks not offered by the manufacturer.

A standard merchandising ploy is to tell the unwary customer that a problem "will be fixed in the next operating system update, which should be out in a couple of weeks." Sometimes this is true, but in far too many cases "a couple of weeks" stretches into months, during which users are forced to work with a product that is bug-ridden or not fully functional. In one or two cases (such as the late, lamented Linn 9000) the company actually went out of business before the updates appeared. Our standard warning applies: Don't buy promises. Make your decision based on what's actually in the store.

On-Board Memory. In general, the more memory a sampler has, the better, because more memory means more seconds of stored sound. An instrument's internal memory is referred to as RAM (random access memory), because the voices have instant access to any sound stored there. Most samplers also have disk drives (see below) for long-term memory storage. The memory figures given may or may not include operating system space in RAM, which will reduce sampling time slightly.

Just to make things confusing, there

Continued on page 35

	FILTER	LFOs	ENVELOPES
AKAI S612*	lowpass	1 global for modulating pitch. Triangle waveform only. Programmable speed, depth, & delay.	Programmable release time only for control of amplitude. V only.
AKAI S700\X7000	12dB/oct. V control of release time.	1 global for modulating pitch. Triangle waveform only. Programmable speed, depth, & delay.	Attack & sustain fixed. Programmable release time. V control of release segment.
AKAI S900*	12dB/oct. V control of cutoff & attack time. Kbd trcking of cutoff. RV. No Q.	1 per voice for modulating pitch. Triangle waveform only. Programmable speed, depth, delay, & response to CAT.	ADSR envelopes for filter cutoff (2.0 operating system) & amplitude. AR envelope for pitch. V control of attack & release.
AKAI S900 w/MARION SYSTEMS MS-9C BOARD	Analog 7-pole, 42dB/oct. lowpass. 92dB dynamic range.	same as S900	same as S900
AKAI S950	12dB/oct. V control of cutoff & attack time. RV. No Q.	1 per voice for modulating pitch. Triangle waveform only. Programmable speed, depth, delay, & response to CAT.	ADSR envelopes for filter cutoff & amplitude. AR envelope for pitch. V control of attack & release; kbd tracking.
AKAI S1000	Digital 18dB/oct. V control of cutoff & envelope attack & release time. RV.	1 per voice for modulating pitch. Triangle waveform only. Programmable speed, depth, delay, & response to CAT.	2 ADSRs w/switchable hold. Control amplitude, filter cutoff, & pitch. V & kbd tracking.
AKAI MPC60	none	none	Attack & fade-down with V control. Programmable rates.
CASIO FZ-1	Analog 4-pole, 24dB/oct. Variable Q. V control of cutoff frequency & Q.	1 per voice w/ sync for modulating pitch, filter cutoff, & amplitude. Sine, sawtooth up & down, triangle, & random waveforms. Programmable delay. Pitch-bend wheel, CAT, & footpedal control.	2 per voice. 8-segment for filter cutoff & amplitude. Positive & negative V control. Programmable amplitude & filter levels & rates, & filter cutoff frequency. Kbd tracking of levels & rates for both envs.
CASIO FZ-10M	Analog 4-pole, 24dB/oct. Variable Q. V control of cutoff frequency & Q.	1 per voice w/ sync for modulating pitch, filter cutoff, & amplitude. Sine, sawtooth up & down, triangle, & random waveforms. Programmable delay. Pitch-bend wheel, CAT, & footpedal control.	2 per voice. 8-segment for filter cutoff & amplitude. Positive & negative V control. Programmable amplitude & filter levels & rates, & filter cutoff frequency. Kbd tracking of levels & rates for both envs.
E-MU EMAX	Analog 4-pole, 24dB/oct. Programmable cutoff & Q. Multiple mod sources.	1 per voice for modulating pitch, amplitude, filter cutoff, & pan position. Triangle waveform only. Programmable rate (w/ random factor) & delay. Programmable by kbd zone.	2 per voice. 5-segment. Positive kbd tracking. V control.
E-MU SP-1200	Six 4-pole, 24dB/oct, 5th-order dynamic filters that can be bypassed.	none	1 per voice. Programmable decay only. No V control or kbd tracking.
E-MU EMULATOR II*	Analog 4-pole, 24dB/oct. Programmable cutoff & Q. Multiple mod sources.	1 per voice for modulating pitch, amplitude, & filter cutoff. Triangle waveform only. Programmable rate (w/ random factor) & delay. Programmable by kbd zone.	2 per voice. 4-stage. V control. Positive kbd tracking.
E-MU EMULATOR THREE	Sixteen 4-pole, 24dB/oct. Variable cutoff & Q. Multiple mod sources.	1 per voice for modulating pitch, amplitude, filter cutoff, & pan position. Sine, triangle, sawtooth, & square waveforms. Programmable rate (w/ random factor) & delay. Programmable by kbd zone.	3 per voice. 5-segment for amplitude, filter cutoff, & auxiliary (pitch, pan position, LFO rate, LFO destination). V independently programmable for each envelope. Kbd tracking.
ENSONIQ EPS	4 low- & highpass digital. 2 per voice in series. Multiple mod sources.	1 per voice. Multiple mod matrix. Sine, triangle, sawtooth, fine triangle, positive-only triangle, & positive-only sine. Random generator. Multiple mod matrix.	6 per voice. 6-segment. Separate envelopes for soft & hard V response (interpolation from one to the other). Multiple mod matrix.
ENSONIQ MIRAGE DSK/ MULTI-SAMPLER*	4-pole, 24dB/oct lowpass. Variable Q. 1 per voice.	1 per voice. Sine waveform only for modulating pitch.	2 per voice. 5-segment for filter & amplitude. V & RV control.
FAIRLIGHT CMI SERIES III	4th-order VCF	1 per voice for modulating pitch & amplitude. Sine, square, & sawtooth up & down. Programmable depth & rate.	1 per sample. 5-segment or drawn on CRT screen. Controls pitch, filter cutoff, & amplitude. V, CAT, & kbd tracking control.
KORG DSS-1*	12 or 24dB/oct lowpass. Programmable Q.	2 per voice, 1 for modulating pitch (either one oscillator or both) & 1 for modulating filter cutoff. Triangle waveform only.	6-segment for controlling filter cutoff & amplitude. V control of attack, decay, & release.
KORG DSM-1	4-pole, 24dB/oct. No Q.	2 per voice, 1 for modulating pitch & 1 for modulating filter cutoff. Sine, triangle, & square waveforms.	6-segment for controlling filter cutoff & amplitude. V control of attack, decay, & release.
KURZWEIL 250 w/SAMPLING OPTION	Analog 9-pole lowpass switch capacitor. 1 per output.	3 per voice w/ sync for modulating amplitude, envs, & pitch. Programmable depth, rate, & delay. Positive-only & negative-only sine, square, sawtooth up & down, asymmetrical, & triangle waveforms.	256-segment for controlling amplitude. V & kbd tracking control.
NEW ENGLAND DIGITAL SYNCLAVIER 3200/9600	none	2 per voice. 24 destinations in modulation matrix. Programmable rate, depth, & delay. 12 waveforms.	5-segment. 128 segments & envelope looping using resynthesis functions. Control each of a sample's partials. V & kbd tracking control.
OBERHEIM-ECC DPX-1	4-pole, 24dB/oct lowpass. Dynamic control depends on emulation scheme.	Same as sampler being emulated (see "Features" column on page 33). Prophet 2000 has only 1 LFO, compared to DPX-1's 8.	Same as sampler being emulated.
ROLAND S-10*/S-220	Fixed, non-real time digital. Low- or highpass. 1 per sample.	1 per voice for modulating pitch. Triangle waveform only. Programmable rate, depth, & delay.	4-segment for controlling amplitude only. V control. No kbd tracking.
ROLAND S-330	Digital 24dB/oct lowpass. 1 per voice. Variable cutoff & Q.	1 per voice for modulating pitch & filter cutoff. Sine & sawtooth up or down waveforms. Programmable depth, rate, & delay.	2 per voice. 8-segment for controlling amplitude & filter cutoff. V & kbd tracking control.
ROLAND S-50	Fixed, non-real time digital. Low- or highpass. 1 per sample.	1 per voice for modulating pitch. Sine waveform only. Programmable depth, rate, & delay.	8-segment for controlling amplitude only. V & kbd tracking control.
ROLAND S-550	Digital 24dB/oct lowpass. 1 per voice. Variable cutoff & Q.	1 per voice for modulating pitch & filter cutoff. Sine & sawtooth up or down waveforms. Programmable depth, rate, & delay.	2 per voice. 8-segment for controlling amplitude & filter cutoff. V & kbd tracking control.
SEQUENTIAL PROPHET 2000*/2002*	4-pole, 24dB/oct lowpass. 1 per voice. V & envelope control of cutoff & Q.	1 for modulating pitch, filter cutoff, & amplitude. Triangle waveform only. Programmable depth & rate.	2 per voice. ADSR for controlling filter cutoff & amplitude. V control. Programmable pedal release.
SEQUENTIAL STUDIO 440	4-pole, 24dB/oct lowpass. 1 per voice. V, env, & kbd control of cutoff & Q.	none	ASRs for filter cutoff, amplitude, & pitch. Sustain either gate time or fixed. V & kbd tracking control.
SEQUENTIAL PROPHET 3000*(YAMAHA P3000*)	4-pole, 24dB/oct lowpass. 1 per voice. V, env, & kbd control of cutoff & Q.	2 per voice w/ sync for modulating pitch, filter cutoff, pan position, other LFOs, & envelopes. Sawtooth up & down, square, & random waveforms.	2 per voice. 5-segment for controlling amplitude & routing in mod matrix. Programmable pedal release. V & kbd tracking control.
SIMMONS SDX	6-pole, 24dB/oct lowpass. 1 per voice. Dynamic control of cutoff & Q.	4 per voice for modulating pitch, amplitude, filter cutoff, pan position, & cross-fade layer balance. Sine, DC, sawtooth up & down, & random waveforms. Controllable with CAT & all MIDI controllers.	6 per voice. 5-segment for controlling filter cutoff & Q, amplitude, pitch, noise generator, & pan position. V, kbd tracking, & pad surface position control.
WAVEFRAME AUDIOFRAME	none	2 per voice for modulating pitch & amplitude. Sine, triangle, square, pulse, random, & sawtooth up & down waveforms. Programmable depth, attack, rate, & delay. Kbd or MIDI controllable.	16-segment for controlling amplitude; 4-segment for controlling pitch. V & kbd tracking control.
YAMAHA TX16W	Real-time digital. 16 types. 1 dynamic axis & 1 fixed axis.	1 global, 1 for filter, 1 per timbre. Global & filter LFOs have variety of waveforms. Timbre LFO has triangle waveform only.	4-segment, 7-parameter. V, B, CAT, & FP control of filter envelope. Kbd tracking control (envelope compression up to a factor of 4).

Key: ADSR, attack-decay-sustain-release; AHDSR, attack-hold-decay-sustain-release; B, breath; CAT, channel aftertouch; CC, continuous controllers; CF, cross-fading; CRT, cathode-ray tube; CS, cross-switching; dB, decibels; DDL, digital delay line; env, envelope; FP, foot pedal; FS, footswitch; HD, hard drive; kbd, keyboard; LFO, low-frequency oscillator; LTC, longitudinal time code; mod, modulation; MW, mod wheel; NA, not applicable; oct, octave; PAT, polyphonic aftertouch; prg, programmable; Q, filter resonance; RAM, random access memory; RV, release velocity; SMPTE, Society of Motion Picture & Television Engineers time code; V, velocity; VCF, voltage-controlled filter; VITC, vertical interval time code.

KEYBOARD LENGTH/ OUTPUT	VOICE RESPONSE	CROSS-FADING & CROSS-SWITCHING	FEATURES	SEQUENCER?
NA	V only	neither	Real-time loop positioning. Manual LFO, filter cutoff, & release-time control.	no
X7000: 5-octave, V & RV	V, RV, & CAT	CF w/version 2.1 only	Retain first 10 samples in memory when turned off. Multi mode for consecutive MIDI-channel voice assignments. X7000 has MIDI local on/off.	no
NA	V, RV, & CAT	CF using V & kbd position CS using V	Positional cross-fades. Input for optional trigger footswitch. HD interface option.	no
same as S900	same as S900	same as S900	Manufacturer reports increase in dynamic range of Akai S900 from 68dB to 85dB. Increase of A/D & D/A conversions from 12-bit to 16-bit. Two modes: 12-bit & 16-bit.	no
NA	V, RV, & CAT	CF using kbd position CS using V	Time stretching--varying playback time without varying pitch. Positional cross-fade. Load samples while playing.	no
NA	V, RV, & CAT	CF using V & kbd position CS using V (overlapping switch zones)	Sample-on coherence. Deflamming of chords. Real-time panning. 320-character graphic LCD. Balanced XLR input jacks. Sample layering.	no
16 V- & CAT-sensitive pads	V; pad pressure used with internal sequencer to repeat sequenced notes	no CF CS using V	99-track sequencer. MIDI note reassignment. SMPTE-MIDI converter. Automated digital mixing w/ graphic display of levels & pan positions. Real-time hi-hat decay & cross-switching using data slider or MIDI.	yes
5-octave, V & CAT	V & CAT	CF using V only CS using V	Graphic display of waveforms & envelopes. Additive synthesis, preset waveforms, hand-drawn waveforms, sample editing for waveforms. Programmable operating system (64K RAM).	no
NA	V & CAT	CF using V only CS using V	Graphic display of waveforms & envelopes. Additive synthesis, preset waveforms, hand-drawn waveforms, sample editing for waveforms. Programmable operating system (64K RAM).	no
5-octave, V only	V & CAT	CF using V, MW, CC, & kbd position CS using V & FS	User-definable MIDI controllers. Microtuning. Auto sample mapping & routing. Arpeggiator; 16-track MIDI sequencer. Multitimbral. 14 velocity curves. Real-time MIDI data filtering.	yes
8 V-sensitive pads	V only	neither	Drum machine--internal response faster than over MIDI. MIDI/SMPTE conversion. Samples can be played in reverse.	yes
5-octave, V only	V & CAT	CF using V & kbd position CS using V	First sampler with synth-type envelope control. MIDI/SMPTE sequencer. Real-time MIDI data filtering.	yes
5-octave, V & CAT	V & CAT	CF using V, MW, CC, & kbd position CS using V & FS	Arpeggiator. Scrub wheel. Real-time loop control. Controller calibration. Legato mode. Auto sample mapping. Analog trigger inputs. Velocity sample start. Microtuning. Real-time data filtering.	yes
5-octave, V, RV, PAT, & CAT	V, RV, CAT, & PAT	CF and CS using mod matrix system	Load samples while playing. Patch-select buttons for real-time control; patch selection using footswitches. Polyphonic aftertouch. Microtuning tables. Stores MIDI sys-ex data for other devices. Plays Mirage disks.	yes
DSK: 5-octave, V only	V, RV, CAT, & PAT	CF using V & MW no CS	Pitch & mod wheel controllers. Velocity & aftertouch response.	yes
76-note, V & CAT	V, CAT, & PAT	CF using V & kbd position CS using V	Synchronization control; video & sequence jogging. Global & local parameter editing. Multitrack hard-disk recording. Multitasking. Audition from hard disk. MIDI/SMPTE sequencers. MIDI processing.	yes
5-octave, V & CAT	V & CAT	no CF CS using V	Oscillator sync. 2 built-in DDLs. Additive synthesis. Data slider can be used to draw waveform. Programmable EQ. Playback at resolutions of 6-, 7-, 8-, 10-, or 12-bit.	no
NA	V & CAT	CF using V only CS using negative V	Additive synthesis. Trigger up to four sounds using footswitch or trigger device.	no
88-note weighted, V, RV, & CAT	V, RV, CAT, & PAT	CF using V & kbd position CS using V	Six sampling modes depending on decay rate of sampled source. Pre-sampling time & threshold programmability. 6-sample layering. Microtuning. Programmable front-panel controls. Real-time sequence editing.	yes
76-note weighted, V, CAT, & PAT	V, CAT, & PAT	CF using V & kbd position CS using V	Optional 2-gigabyte WORM drive for sound database access. 4-, 8-, or 16-track direct-to-disk digital recording & editing for 25 minutes at 50kHz. 200-track automated-recall sequencer.	yes
NA	V, RV, CAT	CF using V & kbd position CS using V	Plays Akai S900, E-mu Emulator II, Ensoniq Mirage, or Sequential Prophet-2000 samples; converts all 4 formats to MIDI sample dump stnd. Converts 8-bit samples to 12-bit. Transpose pitch or sample.	no
S-10: 4-octave, V only	V only	no CF CS using V	Velocity-controlled sample mixing. Dual-oscillator detuning of sample. Arpeggiator with trigger input.	no
NA	V & CAT	CF using V only CS using V	CRT display of waveforms, MIDI data, etc. Outputs assignable by tone & patch. 16-track MIDI sequencer, 8-MIDI-channel multitimbral, 64 samples per channel. Sample cross-mixing. Mouse input.	yes
5-octave, V, CAT	V & CAT	CF using V only CS using V	CRT display of waveforms, MIDI data, etc. 16-track MIDI sequencer, 4-MIDI-channel multitimbral, 32 samples per channel. Sample cross-mixing. Input using digitizer tablet.	yes
NA	V & CAT	CF using V only CS using V	CRT display of waveforms, MIDI data, etc. Outputs assignable by tone & patch, 16-track MIDI sequencer, 8-MIDI-channel multitimbral, 64 samples per channel. Sample cross-mixing. Mouse input.	yes
2000: 5-octave, V only	V & CAT	CF using, V, kbd position, & mod wheel CS using V	Variable sample start point by velocity; negative or positive velocity. Arpeggiator. Stack mode (1, 2, 4, or 8 voices per note). 12 presets.	no
8 V- & CAT-sensitive rubber pads	V & CAT	CF using V & kbd position no CS	SMPTE read & write. 50,000-note, 8-track, 32-MIDI-channel sequencer. Pitch, panning, & volume control from front pannel. Velocity control of sample start point. Audio sample triggering.	yes
NA	V & CAT	CF using V, kbd position, mod wheel, CAT, & MIDI CS using V	Graphic display of waveforms & envelopes. Auto pitch detection (defeatable). Remote-control panel. Auto looping. Macros for filter, amplifier envelope, & mod matrix setup.	no
16 V-sensitive pads (14 positionally sensitive)	V, CAT, & PAT	CF using V & kbd & pad position CS using V	Dynamic range w/ cross-fade & -switch. Sustain pedal control simulates acoustic instrument. Digital mixing. CRT monitor w/ pull-down menus. Integrated system. 64-track sequencer. SMPTE read/write.	yes
NA	V & CAT	CF using V only no CS	Automatable digital mixing. 24-bit stereo digital reverb & processing. Reads & writes LTC & reads VITC. Multitasking; token ring networking. 2 to 32 analog inputs, 4 digital inputs. Audition from HD.	yes
NA	V, CAT, B, MW, & FP	CF using V, CAT, B, MW, FP, & kbd position CS using V	16-note multitimbral. Single triggering. Auto pitch detection. Dynamic digital filters. Programmable velocity curve. CS position displayed graphically. Auto sample mapping. Mono/poly MIDI modes.	no

Continued on page 34

	FRONT PANEL DISPLAY	DISK DRIVE(S)	APPROX. DISK LOAD TIME** (SECONDS)	REPORTED SAMPLING RATES	MONITOR SAMPLING INPUT?	STEREO SAMPLING?	No. OF LOOPS PER SAMPLE	TYPES OF LOOPS
AKAI S612*	none	Optional: external 2.8" QD, 64K bytes per disk	7	4kHz, 8kHz, 16kHz, or 32kHz	yes	no	1	sustain
AKAI S700\X7000	15-char numeric LCD	2.8" QD	112	41 rates from 4kHz to 40kHz	yes	no	1	sustain
AKAI S900*	80-char numeric LCD	3-1/2". 800K words per disk. Optional DMA port for Atari HD.	40	variable from 7,500Hz to 40kHz in steps of 2.5Hz	yes	no	1	sustain, back & forth
AKAI S900 w/MARION SYSTEMS MS-9C BOARD	same as S900	same as S900	same as S900	same as S900	same as S900	same as S900	same as S900	same as S900
AKAI S950	80-char numeric LCD	3-1/2". 1.6 meg words per disk. Optional DMA port for HD, CD, DAT.	40	variable from 7,500Hz to 48kHz in steps of 2.5Hz	yes	no	1	sustain
AKAI S1000	320-char numeric & graphic LCD	3-1/2". 1.6 meg words per disk. HD: Atari, DMA, SCSI.	50	22.05kHz or 44.1kHz	yes	yes	8	sustain, release, programmable time
AKAI MPC60	320-char numeric & graphic LCD	3-1/2". 800K words per disk.	50	fixed at 40kHz	yes	no	none	NA
CASIO FZ-1	128-char numeric & graphic LCD	3-1/2". 1.3 meg bytes per disk. 25-pin parallel port.	45	9kHz, 18kHz, or 36kHz	no	no	8	sustain, release, temporary (time)
CASIO FZ-10M	128-char numeric & graphic LCD	3-1/2". 1.3 meg bytes per disk. 25-pin parallel port.	90	9kHz, 18kHz, or 36kHz	no	no	8	sustain, release, temporary (time)
E-MU EMAX	32-char numeric LCD	3-1/2". 800K words per disk. 20-meg HD optional.	23	10kHz, 16kHz, 20kHz, 28kHz, 31kHz, or 42kHz	no	no	2	sustain, release, loop in release
E-MU SP-1200	32-char numeric LCD	3-1/2". 500K words per disk.	25	fixed at 27kHz	no	no	1	sustain
E-MU EMULATOR II*	32-char numeric LCD	One or two 5-1/4". 650K words per disk. 20-meg HD optional.	36	fixed at 28kHz	yes	no	1	backward, back & forth, loop in release
E-MU EMULATOR THREE	80-char numeric LCD	3-1/2". 800K words per disk. 40-meg HD & SCSI port.	12 (HD)	33kHz or 44.1kHz; other rates from 7kHz to 50kHz available using DSP	no	yes	1	backward, sustain, loop in release
ENSONIQ EPS	22-char numeric fluorescent	3-1/2". 800K words per disk. Optional SCSI port.	31	40 rates from 6.25kHz to 52.1kHz	yes	no	1	sustain, backwards, loop in release, back & forth, real-time loop mod
ENSONIQ MIRAGE DSK/ MULTI-SAMPLER*	2-char numeric LED	3-1/2". 400K words per disk.	7	70 rates from 10kHz to 33kHz; M-S w/ optional filter, 80 rates to 50kHz	yes	no	1	sustain
FAIRLIGHT CMI SERIES III	CRT	380-meg HD. Up to 12 external HDs. Opt. 60-meg tape drive, 800-meg WORM	19 (HD)	continuously variable from 5kHz to 100kHz	yes	yes	1	back & forth, release, sustain, cross-fade, blended
KORG DSS-1*	40-char numeric LCD	3-1/2". 512K words per disk. Optional SCSI port.	40	16kHz, 24kHz, 32kHz, or 48kHz	yes	no	1	back & forth, sustain
KORG DSM-1	48-char numeric LCD	3-1/2". 1 meg words per disk. Reads DSS-1 disks. SCSI optional.	50	16kHz, 24kHz, 32kHz, or 48kHz	yes	no	1	back & forth, sustain
KURZWEIL 250 w/SAMPLING OPTION	48-char numeric LCD	Direct serial port for Macintosh or IBM-PC. RAM cartridge port.	30 to 35 using Mac w/ HD	14 rates from 5kHz to 50kHz	yes	no	1	sustain
NEW ENGLAND DIGITAL SYNCLAVIER 3200/9600	CRT	5-1/4". 320-meg HD. 15-meg cart tape. Optional 2-giga optical.	2.5 per 1 meg (HD)	variable in 1kHz steps from 1kHz to 100kHz	yes	yes	1	sustain & release, backward, cross-fade, butt-splice
OBERHEIM-ECC DPX-1	2-char LED	5-1/4" & 3-1/2". Optional 20-meg HD. Optional CD-ROM port.	15 to 40 (depends on disk)	NA	NA	NA	same as emulated sampler	same as emulated sampler
ROLAND S-10*/S-220	32-char numeric LCD	2.8" QD. 64K words per disk side.	28	15kHz or 30kHz	yes	no	1	sustain, back & forth, reverse
ROLAND S-330	32-char numeric LCD/ CRT output	3-1/2". 1 meg bytes per disk.	35	15kHz or 30kHz	yes	no	1	sustain, back & forth, reverse
ROLAND S-50	32-char numeric LED/ CRT output	3-1/2". 1 meg bytes per disk.	35	15kHz or 30kHz	yes	no	1	sustain, back & forth, reverse
ROLAND S-550	32-char numeric LED/ CRT output	3-1/2". 1 meg bytes per disk. Optional SCSI port for HD & CD-ROM.	70	15kHz or 30kHz	yes	no	1	sustain, back & forth, reverse
SEQUENTIAL PROPHET 2000*/2002*	3-char numeric LED	3-1/2". 1/2 meg bytes per disk.	40	15.625kHz, 31.25kHz, or 42kHz	no	no	2	sustain, release
SEQUENTIAL STUDIO 440	32-char numeric LCD	3-1/2". 800K words per disk. SCSI port.	40	15.625kHz, 31.25kHz, or 42kHz	yes	no	2	sustain, release with alternate
SEQUENTIAL PROPHET 3000*(YAMAHA P3000*)	320-char numeric & graphic LCD	3-1/2". 1.6 meg bytes per disk. SCSI port.	77	16kHz, 32kHz, 44.1kHz, or 48kHz	yes	yes	2	sustain, release
SIMMONS SDX	CRT	3-1/2". 1 or 2 meg per disk. SCSI port. 20- or 70-meg HD opt.	60	11kHz, 22kHz, or 44.1kHz	no	no	1	back & forth, sustain
WAVEFRAME AUDIOFRAME	CRT	5-1/4". 300 meg HD. SCSI. 2.4 gigabyte tape backup. Up to 1.5 gigabyte HD opt.	6 (HD)	44.1kHz	yes	yes	1	sustain, release, cross-fade
YAMAHA TX16W	80-char numeric LCD	3-1/2". 720K bytes per disk. RS-422 port.	80	16kHz, 33kHz, or 50kHz. Stereo at 33kHz only.	yes	yes	1	sustain

Key: CD, compact disc; CD-ROM, compact disc read-only memory; char, character; CRT, cathode-ray tube; DAT, digital audio tape; DC, direct current; DMA, direct memory access; DSP, digital signal processing; giga, billion; HD, hard drive; Hz, Hertz (cycles per second); K, kilo (thousand); LCD, liquid crystal display; LED, light-emitting diode; meg, million; LF, loop finder; M-S, Ensoniq Mirage Multi-Sampler; NA, not applicable; QD, quick disk; RAM, random access memory; SCSI, small computer systems interface; WORM, write-once, read-many optical storage device; **, for full memory load (see "On-Board Memory" column on page 30).

DIGITAL SIGNAL PROCESSING OPTIONS	MIDI SAMPLE DUMP STANDARD
LF.	no
LF.	no
Sample splicing, LF.	no
Same as S900.	no
Sample splicing, time stretch, LF.	no
Sample mixing, gain normalization, LF.	yes
Reverse sample.	yes
Sample mix, cross-mix (splicing), reverse.	no
Sample mix, cross-mix (splicing), reverse.	no
Splice, combine, reverse, gain/attenuate, taper. SE only: pitch change, sample rate conversion, transform multiplication, spectral interpolation.	yes
Reverse.	yes
Splice, combine, reverse, LF.	no
Cut, copy, paste, merge, mix, insert, reverse, digital tuning, taper, gain change, normalization, stereo-to-mono conversion & vice versa, left-right swap, ping pong, auto correlate, auto zero crossing, LF.	yes
Mix, reverse, normalize gain, splice, sample rate conversion, auto zero crossing, similar slope location, invert, cross-fade looping, merge, synthesized loop, bow-tie loop, LF.	no
MASOS (optional software): copy, mix, reverse, fade-in, fade-out, invert.	no
Flanging, sample rate conversion, sample mixing, absolute gain control, fade, digital resampling, LF.	no
Linking & mixing of samples with pitch adjustment (plus or minus 50 cents), sample reverse.	no
Pitch transposition (by octave or semitone), linking & mixing of samples with pitch adjustment (plus or minus 50 cents), sample reverse.	no
Zero crossing search, normalization of gain, pre-emphasis of high-frequency during sampling, de-emphasis on playback (up to 140dB signal-to-noise ratio), cross-fade looping, LF.	no
Sample merging, sample mixing, normalizing gain, phase inversion, time compression & expansion, trimming DC offsets, resynthesis, digital filter construction, LF.	no
NA.	yes
LF with override.	no
Normalize gain, reverse, smoothing, LF. HD5IF software only: DC cut.	no
Normalize gain, reverse, LF.	no
Normalize gain, reverse, smoothing, LF.	no
Sample merging, appending, cross-fade looping, LF.	yes
Normalize gain, sample merging, appending, cross-fade looping (linear & equal power), LF.	no
Normalize gain, cross-fade looping (linear & equal power), LF.	yes
Normalize gain, reverse, linear cross-fade looping, equal-power cross-fade, end sample looping (linear or equal power).	yes
Digital resampling, normalize gain, sample mixing, reverse, cut, copy, paste, scrubbing, LF.	no
Linear cross-fade, normalize gain, wave mix with offset, reverse (whole or mirror), LF.	yes

SAMPLER FEATURES

Continued from page 31

are two ways of stating the amount of memory—in *bytes* and in *words*. A byte is always the same size (eight bits). But the size of a sample word varies. (See "Bit Resolution," below.) On an 8-bit sampler, each word is 8 bits (one byte) long. On a 12-bit sampler, each word is 12 bits (one and a half bytes) long. And on a 16-bit sampler? You guessed it. Each word is two bytes long.

Memory is usually stated in "K" or "meg." "K" is short for "kilo," which is Greek for "thousand." "Meg" means "million." But because computers count in binary arithmetic, "a thousand" may be 1,024. If there are exactly 500 chunks of memory, each of which holds 1,024 bytes, you've got 512,000 bytes, which could be referred to as 512K. Twice as much memory as 512K might be referred to as 1 meg, even though 512 multiplied by 2 would be 1,024K. And ¾ meg could be either 750K or 768K. Is all this perfectly clear?

Next question: How many bytes of memory are there in a 512K sampler? That depends on whether the 512K figure refers to bytes or words. On a 12-bit sampler with 512K, we could be looking at 512K bytes, which translates to about 340K words, or we could be looking at 512K words, which translates to 768K bytes.

The difference is real significant, because it determines how much actual sampling time is available in the machine. If the samples are being played back, for example, at 20,000 words per second, a 512K word 12-bit machine will have 25-plus seconds of memory, while a 512-kilobyte 12-bit machine will have only 17 seconds of memory.

Memory Expansion. A number of sampler manufacturers offer optional memory expansion boards, which can double, quadruple, or more the amount of RAM.

Bit Resolution. This is the ace buzzword that every sampling tyro learns first: Is it an 8-bit, 12-bit, or 16-bit machine? Every time a sampler takes a single sample (a momentary time-slice of the whole sound), it stores the result as a digital *word*. The bit resolution defines the size of the word. Thus an 8-bit sampler stores sounds in the form of long strings of 8-bit words, and so on.

All other things being equal, a sampler

with a higher bit resolution has a better signal-to-noise ratio—as a rule of thumb, about 6dB per bit. Of course, there are lots of other factors that go into sound quality, so a well designed 12-bit machine could outperform a poorly designed 16-bit machine. Some samplers use data companding to get a better signal-to-noise ratio out of fewer bits. Those that don't compand are described as *linear*. The tradeoffs between linear and companding systems are too complicated to explain here, and even engineers don't always agree. Your ears are probably the best guide.

To complicate matters further, a given sampler may not use the same word size throughout its signal path. It may have, for example, 16-bit input and output conversion, but may store the sound in memory in a 12-bit companded format. In this case, marketing types who like to put the best possible face on their product are likely to announce loudly that it's a 16-bit machine, and let you read the fine print to find out that some of the circuitry isn't 16-bit at all. A few will even go so far as to add together the actual bit resolution and the signal-to-noise improvement caused by the analog amplitude envelopes and then announce that their machine has "24-bit signal-to-noise characteristics." Others may report their 24-bit math processors in preference to smaller I/O or memory figures.

When you scrape away the marketing games, the significant factors in bit resolution are (a) the word size in memory, and (b) the word size at the input and output stages. Don't be satisfied with hype about 12-bit, 16-bit, or greater resolution unless you are informed of the actual word size at these points in the signal chain. Our chart lists the three separately, in the form "input/internal/output."

Number Of Voices. A simple matter, you'd think: How many notes can the machine play at a time? Usually eight or 16. Other factors need to be considered, though: When multiple sounds are being layered (or when a single sound is being layered with itself and detuned), a single keystroke usually uses up two or more voices. Instruments like the Ensoniq Mirage and Korg DSS-1, however, have dual oscillators, which means that they can layer two sounds on one keystroke without using up extra voices.

There is another way of performing the same trick (sort of) on a single-oscillator sampler. If the sampler allows you to combine two waveforms into one using a waveform *mix* or *merge* function, you can again play two sounds from one key without using up voices. Mixing is covered in the section of the chart that discusses digital signal processing. In order for two sounds to be mixed in a sampler, they must both be sampled at the correct starting pitch (or digitally pitch-shifted prior to mixing), because pitch-shifting cannot be performed during this type of mixing.

Number Of Sounds & Samples. The simplest samplers (such as the Akai S612) will hold only one sampled sound in memory at a time, but most allow several sounds to be sampled and mapped onto various regions of the keyboard. Our chart gives the total number of samples that can reside in memory at a time, and also the number that can be *played* at a time.

You'll observe that in many cases, more sounds can be played than can exist in memory. What's going on? The answer is that many samplers allow you to use a single sampled waveform a number of times within a single program. Each version of the waveform can be given its own processing parameters, such as envelope and LFO settings. What we've tried to pin down in the "Number Of Sounds" column is the total number of (processed) sounds that you can play from a keyboard or sequencer *without sending the instrument a program change*. Since some instruments allow several programs to be active at a time, it's possible that different manufacturers interpreted our questions about this category in slightly different ways.

Number Of Outputs. Most samplers offer both a mix output (either mono or stereo) that contains the signals from all voices, and a group of individual output jacks. Instruments that have stereo outs may or may not offer programmable stereo pan position or dynamic panning

Continued on page 40

STREET SMART.
BUDGET WISE.

Introducing the DELTEX Keyboard Stand.
Ultimate Quality at an Unbelievable Price.

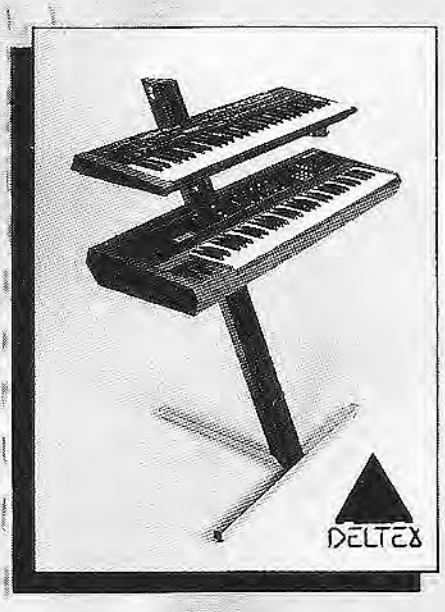

Now every performer can afford the hot new shape in keyboard stands.

With DELTEX, Ultimate Support Systems continues to bring you the same distinctive triangular support column that's become the symbol of rugged strength for stage, studio or home.

Set up and breakdown are fast and easy. DELTEX is lightweight, yet built to shrug off the shocks of the road—a Lifetime Warranty guarantees it.

So check out a DELTEX today. Because now the radical change in keyboard stands is priced to leave you some radical change in your pocket.

Suggested Retail: $99.99
For more information or a dealer near you, contact:
P.O. Box 470
Fort Collins, CO 80522
(303) 493-4488

ULTIMATE SUPPORT SYSTEMS

SAMPLER FEATURES

Continued from page 37

(pan position swept by an envelope or LFO). These features are not tagged in the chart, but they can add a lot to the musicality of a sampler's sound. With some instruments, a voice that is routed to an individual output can be removed from the mix output, which is helpful if you're planning to run it through a separate mixer channel for some individual processing.

On most lower-priced instruments, individual outputs are monaural. That is, each jack receives and passes on the signal from an individual monophonic voice. On a few instruments, the individual outputs are polyphonic and dynamically allocated. In the latter case, you can play, for example, eight-voice string chords from one output jack and then move to another section of the keyboard and immediately get eight-voice brass chords from a different jack, even if the instrument has only eight voices total. Polyphonic, dynamically allocated outputs are especially useful in a small studio where one sampler is being called on to perform several instrumental roles at a time.

On a few instruments, the dynamically allocated outputs are also programmable. By this we mean that the user can set a minimum or maximum number of voices that will be used by a given output, so as to keep the right number of voices available for the various musical parts.

Number Of DACs. This somewhat esoteric-sounding spec relates to the number of outputs and the pitch-shifting technique being employed. A DAC (digital-to-analog converter) is what a sampler uses at its output to convert the numbers it has been shuffling and crunching into analog signals that can be plugged into an amplifier. No DAC, no sound. The simplest samplers have only a single DAC. They may still have multiple voice outputs, if the DAC is multiplexing (time-sharing) among several of them. A better machine will have a DAC for each output. (DACs are not cheap, you see.)

Pitch-Shifting Technique. A CD player will play back digitally recorded sound, but as you may have noticed, you can't speed up or slow down the playback the way you can with a turntable. That's because a CD player has no pitch-shifting capability. A sampler, however, has to be able to play a single sound at many different pitches, depending on what key you play and whether pitch-bend or vibrato is being employed.

There are two ways of getting a digital signal to change its pitch—variable playback rate and interpolation. To understand the difference, imagine a bicycle wheel with a playing card stuck in between the spokes. When we pedal, the wheel spins and the card buzzes. If we want the pitch of the buzzing to go up, we can do either of two things: We can pedal faster, or we can put more spokes around the circumference of the wheel. Pedalling faster corresponds to variable playback rate, while adding spokes is equivalent to interpolation.

Variable playback rate is the method preferred by many designers, because it gives better fidelity when the hardware is in an intermediate price range. The fidelity is better because no matter how quickly or slowly our cyclist pedals, the spokes are still spaced in the same relationship to one another, so their waveform is not distorted. But what happens when we want to add or drop some spokes? How do we decide which ones to drop, or where to add new spokes between existing ones? No matter where we do this, we'll hear some sort of hiccup each time the playing card runs across the spot, unless we do a *lot* of juggling of all the spokes on the wheel. This hiccup is essentially very similar to the digital distortion caused by interpolation.

So why not use variable playback rate in all samplers, if it works better? Well, just as our bicyclist can only pedal at one speed at a time, a DAC can only be playing back at one rate at a time. So what happens when we want to play a four-note chord? If we're using variable playback rate pitch-shifting, we need four DACs. In fact, we need one DAC for each voice of polyphony that the sampler is capable of. And adding DACs, plus their associated clocks and reconstruction filters, adds to the cost of the sampler.

Most high-quality samplers do have one DAC for every voice. Besides minimizing distortion, this arrangement makes it possible for the sampled sound to be passed through individual analog processing circuits (filters, etc.), which makes for more sonic variety.

Less expensive samplers use some type of interpolation for pitch-shifting. With this method, the DAC is always running at the same speed, so if we want to get a three-note chord based on one sampled sound, we have to send the DAC three different versions of that sound. The lower the pitch of playback, the more sample words will be sent to the DAC, and the higher the pitch, the fewer words will be sent. (When multiple voices are being played, an addition process is required as well, so that the DAC receives a digital mix of all the voices.)

If this doesn't make sense, imagine making a tape recording of your voice saying "hello." When you play the tape back at half-speed, two things will happen: The recorded portion of the tape will take twice as long to pass over the playback head, and your voice will sound an octave lower than when it was recorded. Play the tape twice as fast, and again two related things happen: The pitch rises and the sound gets shorter. Now what happens if the "tape" is actually digital data being sent

Continued on page 45

SAMPLER FEATURES

Continued from page 40

to a DAC? If the DAC is running at a constant 20,000 cycles per second, we have to send it 20,000 sample words to create a one-second sound and 40,000 words to create a two-second version of the same sound. Those extra 20,000 words have to come from somewhere.

To lower the pitch, then, we have to add some new sample words in between the words stored in memory. To raise the pitch, we have to drop some of the existing words. This process is called interpolation because new words are interpolated between existing ones. But doesn't interpolation distort the shape of the waveform? That depends on how good a job of interpolation the sampler is doing. With a very fast processor and a very sophisticated software oversampling scheme (see below), you can generate a waveform that is very close to the original, but such processors and software cost money. Unless you're willing to pay a hefty price, the waveform will be more or less distorted. The distortion appears to our ears as noise.

This noise exists only when the sample is played back at some interval above or below its original pitch, and the loudness of the noise will typically change from one pitch to another, because of the changing mathematical relationships between the original waveform, the sampling rate, and the interval of the pitch-shift. Intervals in the equal-tempered scale are not nice whole numbers, so some pretty fancy math is involved in deciding which sample words to drop.

The best way to reduce the noise caused by adding and dropping samples is by *oversampling*. Oversampling is a technique in which the sampler calculates a new, artificial word (or several of them) between each pair of existing words. These words aren't stored in memory, they exist only temporarily during playback. The calculation process is called interpolation. When the sampler has some extra words to play with, it can make better choices about which ones to drop and which ones to send to the DAC. Thus it can produce a smoother (less noisy) output. However, a faster processor must be built into the machine to do the interpolation, which adds significantly to the cost.

Voice Assignment. On a few high-end samplers, you can order a system that has as many voices and as much memory as you need. Most musicians, however, inhabit a world in which each sampler has a fixed number of voices, a fixed number of output jacks, and a (probably somewhat larger) fixed number of sounds residing in its memory. So the question naturally arises, when you play a note on the keyboard, which output jack does it come out of? Many instruments offer their own specialized lists of options in this area. There are too many possibilities for us to cover

them all in detail, but there are a couple of key concepts to arm yourself with before you start pestering the salesperson at your local musitech emporium.

When output jacks are monophonic, they usually exist in a fixed (hardware) relationship to voice channels and/or DACs. But the machine might still assign various sounds to each jack at various points in the music. For example, both open and closed hi-hat sounds could be assigned to the same mono output and removed from the polyphonic mix output. This is called *programmable* voice assignment. When the individual outputs are polyphonic, far more complex types of voice assignment are possible. To save space and avoid endless confusion in our chart, we have indicated only a few basic configurations. For more specific information, consult the manufacturers.

Sampling Rates. As you may have noticed if you read the section on Memory, the faster a sample is being played back, the more memory is being used up per second of sound. Normally, the playback rate for a sound is chosen when it is first recorded. Most samplers offer a selection of two or three different sampling rates, but a few have a wide range of different rates. In order to know which rate to use, you need to know about the overtone characteristics of your input sound.

Here's why: In general, the higher the sampling rate, the higher the frequencies that can be recorded. The upper limit of human hearing is generally put at around 20kHz (20,000 cycles per second), though older people and rock musicians can have significant hearing loss at somewhat lower frequencies—between 10kHz and 15kHz. And many sounds that you might want to sample, such as bass guitar and tom-toms, have no significant overtones above 10kHz. Sounds with weak or nonexistent high overtones can successfully be recorded at a reduced sampling rate, thus increasing the amount of sampling time that can be squeezed from the same amount of memory.

A theorem called the Nyquist theorem tells us that if the sound that is being sampled is to be accurately represented on playback, the sampling frequency must be at least twice as high as the highest overtone in the sound. In other words, for our super-trebly 10kHz bass guitar, we need a sampling rate of only 20kHz. With a 20kHz sampling rate, 10kHz is called the *Nyquist frequency*, because overtones above that frequency must be filtered out, in order to avoid a particularly nasty form of digital noise called aliasing. That's the theory. In practice, a margin of error is built into samplers. The input stage has a lowpass filter that cuts off the overtones somewhat below the Nyquist frequency. To get a reasonable sample with overtones to 10kHz, we would probably want a sampling rate around 24kHz. If the overtones go up to 20kHz, a sampling rate of 44.1kHz (the standard rate for CDs) or even 48kHz would be preferred.

Monitor Input. A minor point of convenience: On some samplers, when you plug a sound to be sampled into their audio input, you can listen to it through their own output. On others, you can't, which can be a hassle unless you're feeding the input through a send on your mixing console.

Stereo Sampling. A few of the newer samplers are offering true stereo sampling. This can be useful for making almost anything sound bigger, and is indispensable if you're planning to sample bits of a performance off of one section of tape and fly them in to a different section by triggering the sampler. Of course, sampling in stereo uses up twice as much memory.

Loops. The art of looping samples, so as to make a short sound sustain forever, is one of the most important (and most intimidating) aspects of using a sampler. Just about all samplers will let you set the length of the loop, and many offer an automatic loop finder—a software procedure that attempts to find smooth loop points for you. Some allow cross-fade loops to be constructed.

Many samplers offer other looping options more complex than single loops. You may be able to set two or more separate loops within a single sound, and define how the machine will play through them. For example, a *sustain loop* is a loop that cycles over and over until you lift the key, at which time the rest of the sample starts playing back.

Some samplers allow back-and-forth looping, which can be useful for smoothing out a loop and making it longer. At least one machine allows the loop length to be altered in real time by the mod wheel or other performance controller. Another lets you define the total time length during which the loop will cycle; in this case, a lower note will play the loop fewer times than a higher note, because the higher note will get through the loop faster.

Digital Signal Processing. The higher-priced samplers offer many ways to alter a raw sample once it is stored in memory. Even the simplest samplers will allow you to truncate (shorten) a sample, while on better machines you may be able to splice two samples together, mix them, cross-fade from one into another, or normalize their gain (increase the signal level to the maximum allowed by the machine).

If you've ever tried to mix two sounds together, you may have found that one of them was at the wrong pitch, and that there was no way to retune it, because pitch-shifting was strictly a playback function. Being able to pitch-shift sounds in memory is one of the DSP options; another is digital EQ. Digital enveloping, delay functions, and stereo imaging edits are also found on a few samplers.

If the sampler of your choice doesn't offer much in the way of DSP, don't despair. A number of computer programs are available that will receive and transmit sound files using the MIDI Sample Dump Standard. These programs provide some rather esoteric DSP functions, and they will work with those samplers designated in the "MIDI Sample Dump" column in the chart, as well as those that are specifically addressed by particular programs. In the course of running the audio tests for this issue, however, we found a number of small but thorny problems in SDS implementations on various instruments. So before you rush out and buy a Macintosh and an expensive piece of software to fold, staple, and mutilate your sampled antelope belches, talk to some people who are currently using the same equipment to find out whether they've encountered unsolved problems.

Display. The most basic configuration is a two-digit LED display. Many samplers offer multi-segment LCDs, which can display a whole batch of parameters at once in numerical form. A few have large enough LCDs that some elements, such as waveforms and envelopes, can be displayed graphically. A few others have outputs to drive a computer-type monitor screen.

Disk Drive. A few samplers (such as the Kurzweil 250) require external hardware in order to load new sounds into memory, but most have their own disk drives. Three sizes of disks are in common usage—2.8" (also called quick disks), 3½", and 5¼". 3½" disk drives come in single-sided and double-sided formats; the double-sided ones obviously allow twice as much data to be stored on a single disk. Double-sided drives can be either double-density or high-density, and here again, high-density disks can store twice as much data.

Some samplers are available with an optional built-in hard drive, which allows a great deal of sound to be stored with no messy disk-swapping, and loaded very quickly.

The figures in the chart under "Disk Load Time" indicate how long it takes to load the instrument's entire basic memory configuration from its standard disk drive. Load times with optional hard disks are not given, but in the case of the Emulator Three the hard disk is standard rather than optional, so we used it to figure the disk load time.

LFO, Filter, & Envelopes. These are the main *analog processing* features. A sampler may be handling its audio entirely with digital circuitry, but the term 'analog' may be used anyway, to indicate that these features are similar to those found on older analog synthesizers.

The way in which analog parameters are assigned to raw sampled sounds is determined by the sampler's *architecture*. On most instruments, a single sample can be assigned to several regions of the keyboard, and each region can be given its own envelope, filter, LFO, and other characteristics. The architecture through which this setup is programmed, and the names used for the various elements. are not standardized. So it can be a little tricky to understand what's going on in a new instrument.

Most samplers use digital LFOs for vibrato, because all pitch-shifting is being done in the digital domain. (See below for a close look at the varieties of pitch-shifting technology.) Usually, each voice/sound/region will have its own LFO, but there may be circumstances, such as several voices playing the same sound, in which voices will share an LFO.

Some samplers, but not all of them, have analog-style lowpass filters. These are used for making the sound brighter or more muted. (Samplers also have lowpass filters at their sampling inputs and smoothing filters at their outputs, but these may be invisible to the user, and are never programmable as part of a preset.) Very few samplers offer filter *resonance*, however. Resonance is found mostly on analog-style synthesizers.

Most samplers, except for a few drum machine types, have amplitude envelopes, which can be used to shape the loudness of the note—by muting part of the attack or making a quicker decay, for example. If a sampler has a lowpass filter, it may also offer a filter envelope, which allows the filter to be swept dynamically during the course of the note.

The most common envelopes are ADSR (attack-decay-sustain-release) types. Some samplers have more complex multistage envelopes, or even envelopes that permit several stages to be looped. Multistage envelopes are useful for a few types of sounds, but are not necessary at all for many sounds.

Voice Response. This column explains what types of incoming MIDI data the sampler responds to. The column labelled "Keyboard Length/Output" explains what types of performance data are generated on the sampler's own keyboard. Some keyboards transmit only velocity, some have monophonic (channel pressure) aftertouch, and others have polyphonic (key pressure) aftertouch. A few also transmit release velocity.

MIDI control messages can be used to modulate lots of different aspects of the sound. Some instruments, for example, provide aftertouch but only allow it to modulate vibrato depth. On others, aftertouch may be used to open up the filter, bend the pitch, or change the loudnesses of various sounds in a layer. Especially if you're planning to use a sampler with a wind controller, you need to look closely at how it responds to controller information. Our chart lists only the types of information that can be responded to, not the destinations to which the control can be routed.

Cross-Switching & Cross-Fading. Most samplers have setups that allow more than one sound to be played from a single key. There are several ways of doing this. The simplest is *layering*, in which two or more voices are used with every keystroke. Some instruments allow more complex types of layering called velocity switching and velocity zoning. In a velocity switch, only one sound will be played with each keystroke, but which sound depends on the key velocity. For example, a soft snare hit might be assigned to low velocities (such as 1-72) and a hard snare hit to high velocities (73-127).

With velocity zoning, each of the sounds in a layer can be given an independent zone of velocities, and will play whenever a note within this velocity range is played. This setup allows velocity zones to overlap. For example, the soft snare sound could play at velocities 1-127 and the loud one (or some other reinforcing sound) could also play at velocities 92-127.

A still more flexible scheme is *velocity cross-fading*. In this case, when two velocity zones overlap, one sound will fade out as the velocity increases, and the other will fade in.

Some samplers offer *positional cross-fading*. This allows the transitions between one keyboard zone and another to be smoothed out by playing both sounds in the overlapping region and adjusting the relative loudnesses of the two sounds based on what note is being played. Note that both velocity cross-fading and positional cross-fading use up two voices per note, since two independent sounds are being played.

SAMPLERS LAID BARE

HOW WE SET UP OUR LAB TESTS & WHAT WE LEARNED

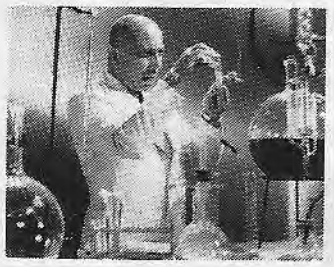

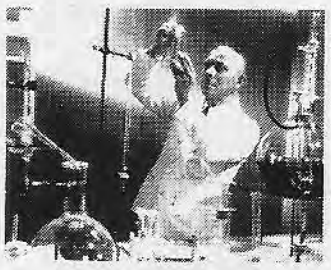

BY TED GREENWALD

EVEN UNDER THE BEST HIGH-RES, high-rate, brick-wall circumstances, life as a digital sample must be a disorienting thing. After all, a sample file is a sound without any sound in it. It may be encoded on DAT, CD, ROM, RAM, or home video. It may shuttle endlessly between floppy disk, computer network, and sampler, only to be filtered, truncated, looped, and subjected to all manner of mathematical insult, from data compression to sample-rate conversion. After sampling, what was once clear, ringing *sound* in all of its mysterious splendor is transmuted into a mere cipher of ones and zeroes, in many ways indistinguishable from MIDI, a FAX transmission, or NORAD's command to kick off World War III.

As musicians during the dawn of the digital age, we must regard these circumstances as facts of life. But—as musicians—we also care about our sounds and what happens to them. If they must be whisked off to digital hell, we owe it to them, at least, to know what they're going through. Besides which, we're just plain curious.

With the help of as many experts as we could enlist, we set out to put a stethoscope to the heart of virtually every professional sampler on the market, from the competitive price point to the state of the art. What began as a few routine checkups eventually ballooned into a crash course in digital brain surgery, and 13 weeks later we're ready to take off the surgical mask, peel away the bandages, and present our Frankenstein to the world.

In the pages that follow, we'll evaluate the implications of such an undertaking, explain the tests and how to interpret their results, and highlight some of the stranger issues that came to light. Interspersed amid the tangle of digital audio theory and practice, you'll find both side-by-side graphic comparisons and individual profiles detailing the strengths and weaknesses of each instrument.

So relax. Take three deep breaths. And try not to jump too much when we tap you on the knee with a Kurzweil 250. We promise that, when we finally put your sampler back together, you *will* play the violin again.

DR. KEYBOARD'S GOOD SPECS

SOUND IS AN OBJECTIVE PHENOMENON. *Good* sound, being a subset of sound in general, is also objective. But, obviously, good sound includes a crucial subjective component. Specification measurements, better known as specs, are an attempt to objectify the subjective aspects of good sound—that is, to point a finger at the things that differentiate a "good" sound from a "bad" one, and then measure them. As we know from the history of art, philosophy, and the great religions, trying to reconcile the objective and the subjective is always a tricky business. Thus, it's a truism among the audio community that specs don't tell the whole story.

At heart, a sampler is an audio recording device, and how good it sounds depends in large part on the quality of its recording and playback circuitry. Audio technicians have come to regard the measurable aspects of a recording device's performance as *frequency response*, *dynamic range*, *distortion*, and *transient response*. Frequency response re-

PHOTOS BY BEITMANN ARCHIVE

fers to the range of frequencies over which the device can record and play back, dynamic range refers to range of amplitudes over which the device can record and reproduce, distortion refers to the accuracy with which these ranges are handled, and transient response involves the speed with which the circuitry reacts to changes in the above.

Still, the relationship of these characteristics to good sound isn't a simple one. Engineers point out that a particular spec that looks good on paper can end up sounding bad, simply because the human ear is sensitive to imperfections in that area. On the other hand, there are areas in which the ear is much more forgiving, in which poor performance won't sound particularly bad. For instance, even-order harmonic distortion (consisting of every second harmonic, beginning with the fundamental) is said to be tolerable to the ear even at high levels. Odd-order harmonic distortion in equal amounts doesn't sound as sweet.

Furthermore, expecting a sampler to be a perfect digital recorder is somewhat impractical. A manufacturer may be able to bring a sampler's cost down considerably by compromising a bit on fidelity, allowing for the addition of features that will enhance its usefulness. Also, a sampler is called upon to do things that a dedicated recording device never has to do. These include playing looped segments of recorded sound and playing several pitch-transposed versions of the same recording simultaneously. A sampler is as much a musical instrument as it is a recording device, so to evaluate it solely as a recording device is to ignore other dimensions of its functioning.

If specs don't tell the whole story, what does? The usual answer is your ears. But no one has the opportunity to listen to every instrument, and, given stress, fatigue, emotions, and so forth, you can't always be sure that your ears are hearing everything there is to be heard. So it makes sense to have some cold, hard facts available to supplement what your ears tell you. Specs may not tell the whole story, but they do paint a picture which is framed in the objective language of science rather than the subjective terms of human perception, and that can be useful when you're not sure what your ears are telling you.

TESTING FREQUENCY RESPONSE

ANY AUDIO DEVICE, WHETHER IT'S A speaker, a synthesizer's lowpass filter, a preamp, or whatever, will perform more or less effectively depending on what range of frequencies it's called upon to handle. Ideally, frequencies presented at a sampler's input should come out the other end at the same level. This is what we mean by frequency response. Human hearing generally spans the range between 20Hz and 20kHz (known as the audio range), so we're most concerned with what goes on in that area.

We tested frequency response by recording and playing back white noise. White noise, averaged over time, consists of an equal distribution of energy throughout the audio spectrum. A spectral analysis of a sampler's reproduction of noise, which we performed using a Hewlett-Packard 3561A Dynamic Signal Analyzer, provides a strikingly graphic representation of its frequency response curve. We'll go further into how to read spectrum plots in a moment. For the time being, have a look at the one reproduced below, which depicts the noise, coming directly from its source, that we sampled into each instrument. (The plot's vertical axis represents level in dBV; the horizontal axis is frequency in Hz.)

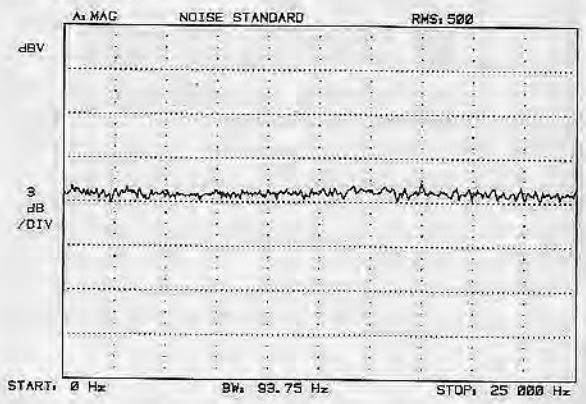

The flatness of the curve indicates that all of the energy in the white noise is represented. Any dips or peaks would indicate inaccuracies—that is, a failure to reproduce the energy originally in the white noise. (Some readers may be accustomed to seeing frequency response plots with a logarithmic, rather than a linear, frequency scale. In the context of this article, we feel that it's less confusing to use linear plots, since our distortion tests are plotted the same way. The trade-off is that a linear plot gives less detail about low-frequency response.)

A sampler's frequency response is bound to differ from that of our noise source in that it will roll off in the upper range. Due to a digital phenomenon called *aliasing*, in which high frequencies are transposed into rather ugly low-frequency components, the frequency response of a sampler's input circuitry should cut off sharply somewhere below half the sampling rate. (This function is performed by the *anti-aliasing filter*.) Likewise, the mathematics of sampling dictate that components related to the sum and difference of the sampling rate and the frequencies sampled (called *imaging*) will appear in a sampler's output. Therefore, the output circuitry should also roll off below half the sampling rate. (This is the job of the *reconstruction filter*, also called the *anti-imaging filter*.)

The range of an instrument's frequency response before it rolls off is called the instrument's *bandwidth*. (In the instrument profiles beginning on page 55, bandwidth is considered the range between 0Hz and the frequency at which the level has decreased by 6dB from the level at 1KHz. We consider this a liberal definition.) As you can see, bandwidth is related to sampling rate. Generally speaking, the higher the sampling rate, the wider the bandwidth, the more high-frequency material is preserved, and the better we can reasonably expect the sampler to sound. Likewise, the sharper the cutoff at the upper limit of the sampler's bandwidth, the lower the possibility that aliasing will occur, and the better the sampler should sound.

Let's take a look at our noise source after it's been sampled and played back through an Akai S900.

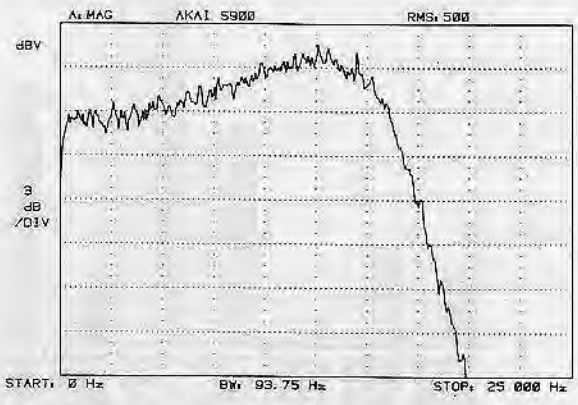

It's clear that the S900 has anything but a flat frequency response. Its most unusual feature is a pronounced emphasis of the range between 7.5kHz and the rolloff point, at

Acknowledgments:

Test Design: *Dave Clementson (Digidesign).*

Additional Test Design & Digital Audio Consultation: *Marcus Ryle (Fast Forward Designs), Tom Oberheim (Marion Systems), Dave Rossum (E-mu), Mark Jeffery (Digidesign), Rick Downs (Burr-Brown), Bob Chidlaw (Kurzweil), Bill Mauchly (Ensoniq), Bill McCutcheon (Ensoniq), Brent Hurtig (GPI Publications), Jim Letts (Oberheim/ ECC).*

Test Equipment: *E-mu Systems, Digidesign, Test Equipment Corp., Rent-A-Computer.*

Special Thanks: *Peter Gotcher and the entire Digidesign staff, Paul Wiffen, Chris Now (Digidesign), Evan Brooks (Digidesign), Bill Gardener (Kurzweil), Tom Metcalf (Ensoniq), Jim Lindop (Simmons), Bob Moog (Kurzweil), Mark Suruta and Marty Frasu (Roland), Ed Alstrom (Casio), Greg Grattan (Akai/IMC), Jack Hotop (Korg USA), Charles Fielding (Yamaha), Matt Tietjen (Man-With-A-Van), Peter Buffett (Independent Sound), Michael Macrone (Statistical Concepts).*

around 16kHz.

If flatter is better, does the graph indicate that the S900's frequency response is bad? If the S900 were a CD player or a tape deck, we'd say yes. However, as a musical instrument, the S900 is prized for its punchy sound, and some of that punch certainly comes from its oddball high-frequency emphasis. If your sampler applications require accuracy—such as flying a single chorus into several parts of a song—the S900 won't do the job as well as a machine with a flatter frequency response. However, if you want a bright, punchy sound for your percussion samples, you're more likely to get it from an S900 than from a sampler whose high end is flat right up to 16kHz.

MEASURING DISTORTION

THE FREQUENCY RESPONSE TEST PAINTS a sampler's capabilities in a single broad stroke. It indicates how well the instrument is capable of recording and reproducing frequency-dependent energy throughout the audio spectrmum, but it says nothing about how precisely it handles that energy. Imprecision in this regard is called distortion, which may be thought of as impurities in a recorded and/or reproduced sound.

In order to measure audio impurities, we start with an audio signal which we know is pure: a sine wave. A perfect sine wave consists of a single, isolated frequency, which we hear as a fundamental tone. Distortion introduced into our perfect sine wave appears as additional frequency components. A measurement device—in this case, a Tektronix AA 501 Distortion Analyzer—can filter out the fundamental, measure what's left, and give us a percentage representing the amount of distortion. This percentage is called THD+n, for *total harmonic distortion plus noise* (although it includes inharmonic distortion).

Distortion comes in three flavors: harmonic, inharmonic, and plain old noise. Our trusty Tektronix doesn't know one from the other. Your ears, however, do. Given that, in a recording device, all distortion is undesirable, harmonic distortion is good distortion. Harmonic distortion appears in whole-number multiples of the fundamental frequency (and sometimes as whole-number divisions of the fundamental, called *subharmonics*). Because it's related to the fundamental in ways that accord with the harmonic series, harmonic distortion sounds like a change in timbre. Inharmonic distortion, on the other hand, bears no simple mathematical relationship to the fundamental. It simply sounds as though it isn't supposed to be there. This is bad distortion. And then there's noise; as we all know, no noise is good noise.

The point is that THD+n indicates how much distortion is present, but not how benign or nefarious it is. Therefore, our THD+n measurements should be regarded only as general indicators, not as an absolute basis of comparison.

A comparison of average distortion. The red bar shows the amount of distortion introduced when a sine wave is sampled through the instrument's input stage and played back through its output stage. (There is no input-plus-output distortion figure for the Oberheim DPX-1, which is a playback-only machine.) The amount shown is an average of the weighted THD+n measured at three different frequencies—97Hz, 997Hz, and 10,007Hz. The yellow bar shows the average amount of distortion introduced by the output stage alone.

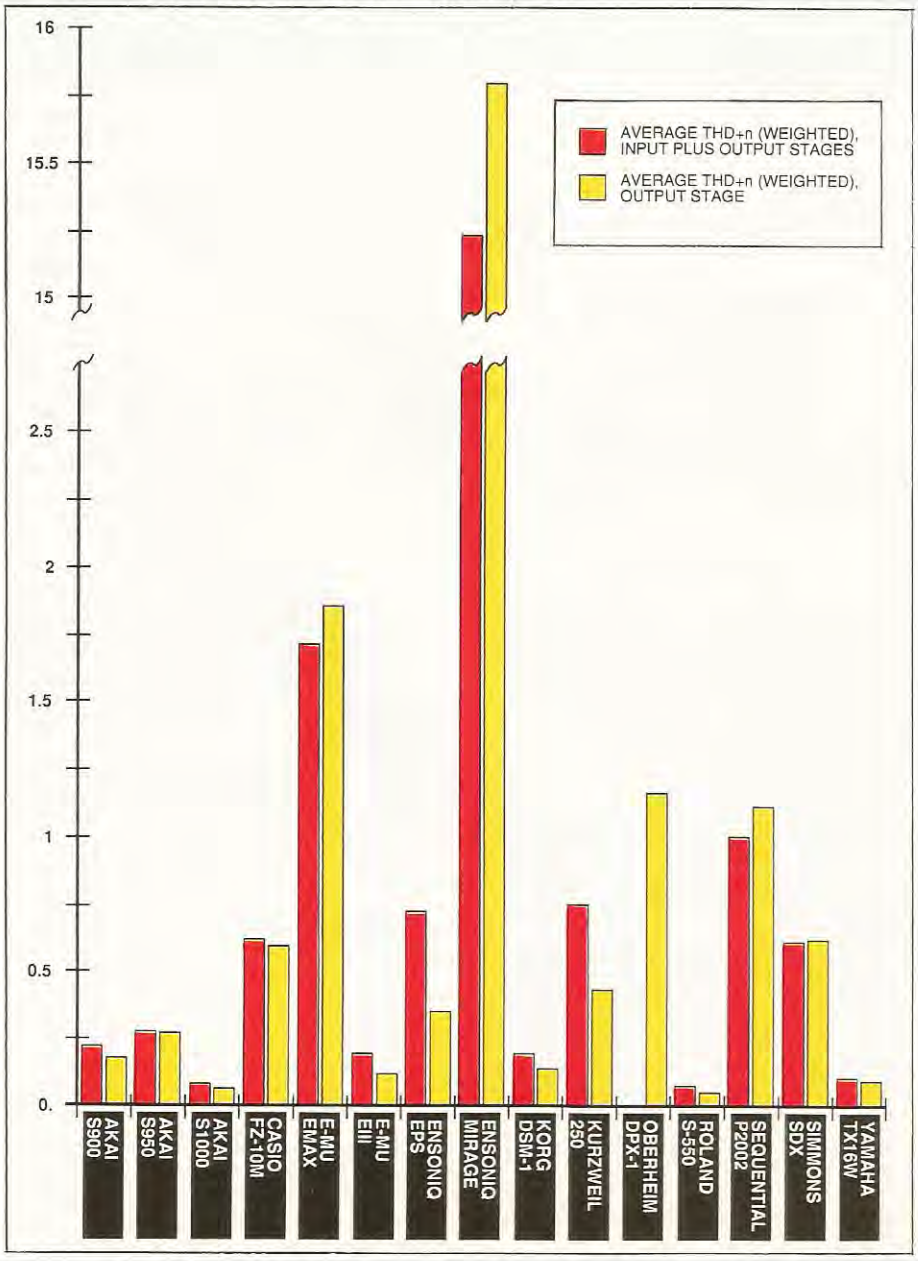

A-WEIGHTING & BAND-LIMITING

ANOTHER OBSERVATION ABOUT good and bad distortion: Harmonic, inharmonic, or what-have-you, distortion sounds worse in the middle portion of the audio range than in the lower and upper. That's because the ear is more sensitive in this range. Audio engineers have used this phenomenon as a basis for the A-weighting curve, a mathematical function designed to compensate for the difference between what we hear and the frequency components actually present in a sound. Since, on average, human hearing rolls off above 10kHz, A-weighting makes distortion in the range above that frequency count for less in the overall THD+n percentage. The ear is most sensitive from around 3kHz to around 7kHz, so A-weighting makes distortion in that range count for more. Here's what the A-weighting curve looks like (plotted against an exponential frequency scale):

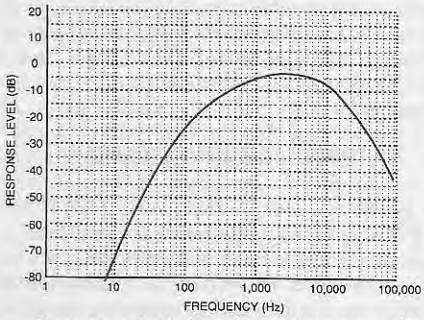

Bear in mind that distortion can also occur far above the audio range. In other words, your axe may sound great to you, but it can be an aural nightmare for the family

dog.

Should human beings care what goes on above the range of their hearing? Well, yes and no. Yes, because other components in the signal chain (including amplifiers, tape decks, and signal processors) may respond to supersonic frequencies in such a way that they become expressed in the audio range. Yes, also, if you ever re-sample the output of an instrument that's spitting out supersonic frequencies. If the device doing the resampling has a wide enough bandwidth, the supersonic frequencies will be recorded properly, and you'll hear them loud and clear as they're pitch-transposed down into the audio range during playback. If the device doing the resampling (which includes feeding it into a digital signal processor or a digital recorder) has a narrower bandwidth but a poor input filter, the supersonics will appear as aliasing.

On the other hand, there seems to be some disagreement over how common the circumstances described above actually are. Some of the experts with whom we consulted felt that supersonic garbage is worth avoiding, while others assured us that in the hi-fi world of the '80s, superonic frequencies are likely to stay out of earshot.

All of which explains why there are two sets of numbers in our THD+n chart on page 55. The lower number represents the amount of distortion actually present in the signal being measured. The upper number differs in two important ways: First, it has been A-weighted. Second, the A-weighted figure has been band-limited at 30kHz (that is, a lowpass filter has knocked out everything above that frequency), further compensating for the way the distortion is likely to sound. We believe this provides a more meaningful indication, in musical terms, of an instrument's distortion characteristics. For this reason, we relied exclusively on the A-weighted and band-limited figures in making up the graphs on pages 52, 53, and 54. (For the remainder of this article, "weighted" should be taken to mean "A-weighted and band-limited.")

The significance of the adjusted THD+n numbers is demonstrated most dramatically by the Ensoniq EPS. A 997Hz sine wave sampled and played back through the EPS at the instrument's 48.1kHz sampling rate picks up nearly 4.2% distortion along the way. Pretty appalling, huh? If you extend the standard 25kHz range of our distortion plot to 100kHz, it's plain to see that much of that

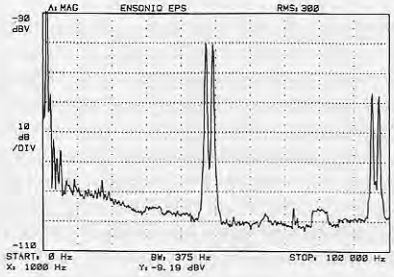

percentage consists of two pairs of spikes centered around 48kHz and 96kHz, well above the range of human hearing.

Weighted, the THD+n measurement

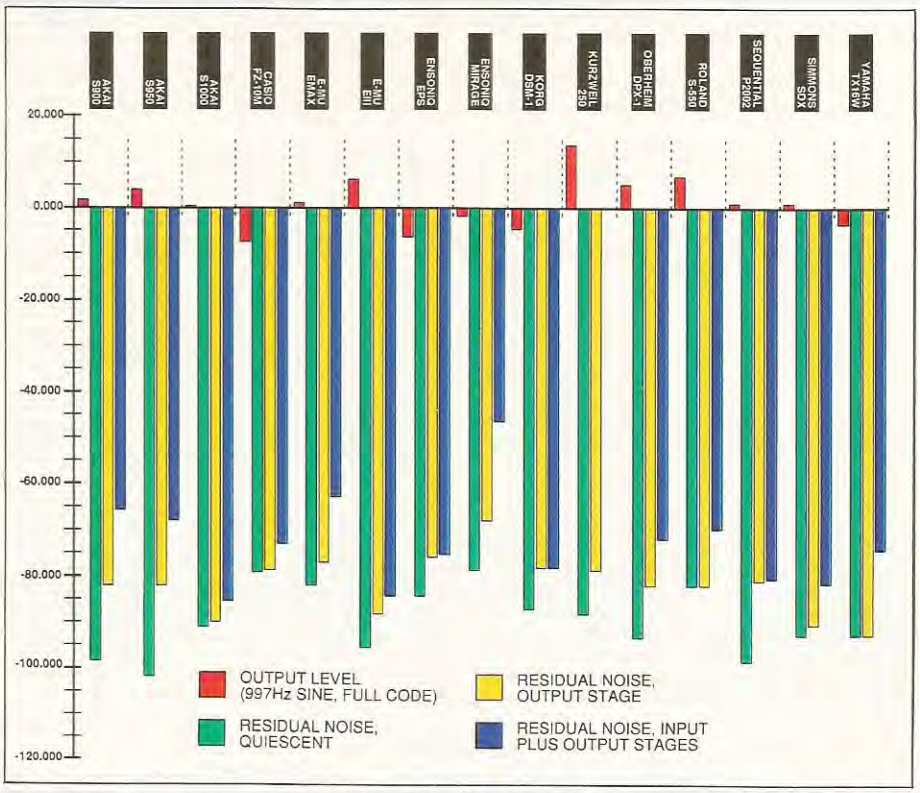

A comparison of dynamic range measurements. The red bar indicates the output level, in dBm, of a single voice playing back a full-code 997Hz sine wave through the instrument's mono mix output, with the volume level full up. The three noise floor measurements are relative to this level, which is considered 0dB. (The higher the output level is, the lower the instrument will normally be set at the mixer, which in turn will reduce the amount of extraneous system noise.) The quiescent noise level is the amount of noise present when the instrument is idle. The output-stage residual noise level is the amount of noise present when the instrument plays back a sample file consisting of zero-amplitude words. This represents the noise generated by the output stage alone. The input-plus-output-stage residual noise level is the amount of noise when the instrument plays back sampled silence. In most cases, the silent sample was taken with nothing plugged into the instrument's input jack, and the input level control at two o'clock.

drops to a more reasonable .47%. This is a much more realistic indication of what the EPS's distortion at 997Hz actually sounds like.

THE PLOT THICKENS

As you can see, spectrum plots make it easier to evaluate our THD+n measurements, allowing us to see the composition of the distortion—its level, frequency range, and harmonic relationship to the fundamental. The H-P analyzer generates a spectrum plot by dividing the frequency spectrum into a bunch of very narrow bands, called *frequency bins*, and then measuring the amount of energy present within each band. The result is a kind of a bar graph of the frequency content of the input signal. Don't confuse this kind of representation of sound with an oscilloscope-style representation. Sound exists in three dimensions: time, amplitude, and frequency. An oscilloscope or a graphic waveform display on a computer plots amplitude against time. A spectral analyzer plots amplitude against frequency.

A spectrum plot of an ideal sine wave would show a single peak, at the sine wave's frequency, rising from the noise floor of the electrical environment (produced by the signal generator, the spectrum analyzer, and the house current) to a level representing the sine wave's amplitude. The closest thing we could find to a perfect sine wave, at least one whose frequency and amplitude could be precisely controlled, comes from a Hewlett-Packard 3336B Synthesizer/Level Generator. A plot of a 997Hz tone at nearly 0dBV, generated by the H-P 3336B in the electrical environment of *Keyboard*'s offices, looks like this:

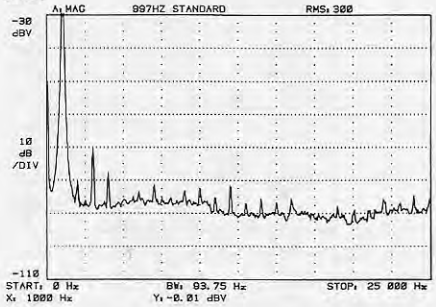

The sine wave's peak, on the left-hand side, extends beyond the graph's upper border. The distortion components rising from the noise floor to its right are harmonic (you can tell by the even spacing between them). Although it's a bit disconcerting to see distortion so clearly in our test waveform, the THD+n measures .0544% (weighted), which is quite low.

This 997Hz sine wave was one of the three test signals we used to measure the combined distortion induced by sampling a sound and playing it back. Plots of the other two test tones, 97Hz and 10,007Hz, are reproduced below. You'll find them useful for

comparison when you look at the profiles of individual instruments, which begin on page 56.

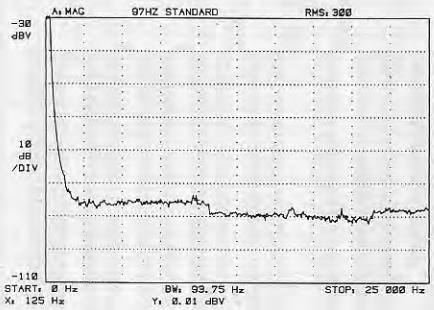

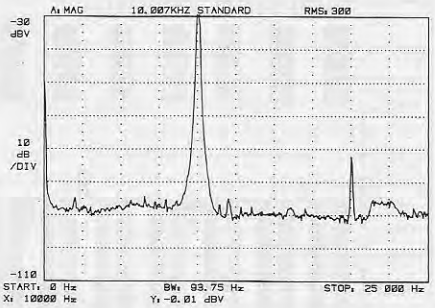

HOW TO READ SPECTRUM PLOTS

THE HORIZONTAL SCALE OF A SPECtrum plot represents frequency in Hertz. In the plots presented here, it extends from 0Hz to 25kHz—slightly beyond the audio range in both directions. Each horizontal division represents 2.5kHz. (They're easier to read if you count every two divisions as 5kHz.)

The vertical scale (level) is calibrated in decibels, or dBV, with each division representing 10dBV. The plots of our test signals show the range between -30dBV and -110dBV. Be sure to note this scale whenever you look at the plots reproduced in these pages. We've tried to keep it consistent, but sometimes we had to move it up or down in order to keep a given signal's peaks and troughs within the borders of the graph.

In the lower left-hand corner you'll notice two numbers designated "X" and "Y." These are the coordinates of the fundamental frequency's peak, which extends above the graph's boundary. In this case of the 997Hz signal, X equals 1kHz (1000Hz) because the analyzer's resolution isn't fine enough to register 997Hz. Likewise, X in the 97Hz plot reads 125Hz, and in the 10,007Hz plot, 10kHz. In the plots above, the Y coordinate shows that the amplitude of the fundamental is nearly 0dBV. This number varies from instrument to instrument and from frequency to frequency according to the strength of the instrument's output level and the curve of its frequency response.

Note: *It's crucial to note the Y-coordinate, the amplitude of the fundamental frequency, whenever you try to interpret one of the distortion test plots. That's the only way to get a realistic idea of the strength of the noise and distortion.* For instance, if the fundamental is at +10dBV and the loudest harmonic at -40dBV, the distortion is actually 50dBV lower than the fundamental. If the fundamental is at -25dBV—much lower in level but still off the graph—a distortion component appearing at -40dBV would be only 15dB down, which is a much more severe problem.

A few incidental aspects of the spectrum plots presented here: The number labelled "RMS" indicates the number of times the analyzer looked at the input signal in order to arrive at an average representation. It equals 300 for almost all distortion plots, and 500 for all frequency response plots. "A: MAG" and the number associated with "BW:" refer to more arcane functions, beyond the scope of this article, used in the analyzer's computations.

HERTZ SO BAD

TEST FREQUENCIES OF 97Hz, 997Hz, AND 10,007Hz? Why such messy numbers? Why not friendly, civilized figures like 100Hz, 1kHz, and 10kHz? Because, digital circuitry being based on mathematics, the numbers involved make a difference. An analog audio device will behave the same way whether you feed it 997Hz or 1kHz, barring design flaws which may cause inconsistent performance. With digital audio, however, performance varies with the frequency being recorded or played back.

Imagine you're an analog-to-digital converter. You're being clocked by a slave-driving microprocessor at 40kHz, which

A comparison of the average amount of distortion added by pitch transposition. The yellow markers represent the average difference between the THD+n (weighted) of a full-code 997Hz sine wave transposed by a half-step, a minor sixth, an octave, and an augmented 11th in either direction and the THD+n (weighted) when the same sample is played back untransposed. Red markers represent the THD+n (weighted) when the sample is played back untransposed, and are provided for comparison. Performance is closest to ideal where the yellow marker falls nearest the horizontal axis, representing no change introduced by transposition. Performance is least ideal where the yellow marker is above the red marker. "Negative" pitch-shift distortion levels are associated with high untransposed THD+n measurements, and may result from distortion components being transposed beyond the ranges of the weighting and band-limiting curves.

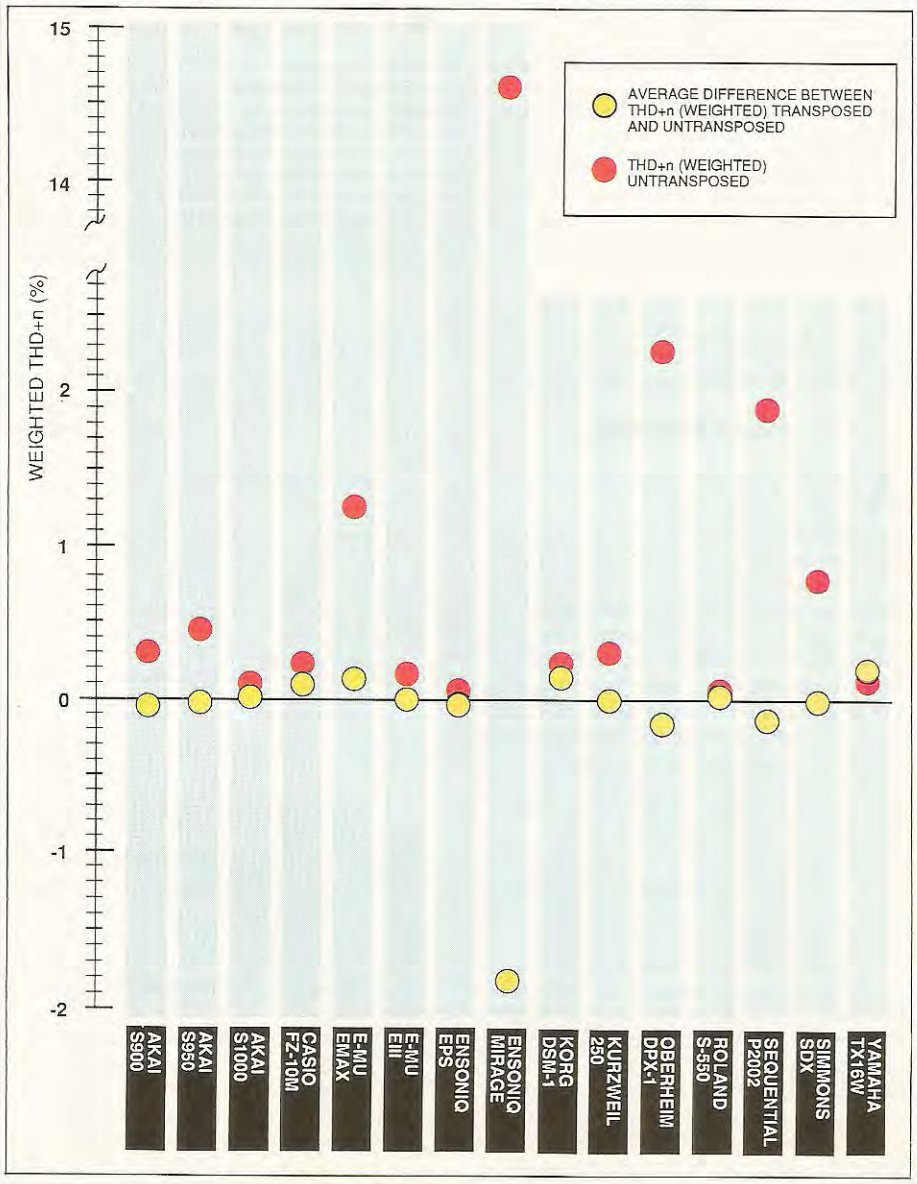

	SAMPLING RATE (Hz)			DYNAMIC RANGE (dB)				INPUT + OUTPUT STAGE DISTORTION (%)			OUTPUT STAGE DISTORTION (%)			PITCH-SHIFT DISTORTION (%)							
														UP				DOWN			
	REPORTED IN MANUAL	EXPERIMENTALLY DERIVED		QUIESCENT NOISE FLOOR	SILENT SAMPLE OUTPUT STAGE ONLY	SILENT SAMPLE INPUT+OUTPUT STAGES	OUTPUT LEVEL 997Hz FULL-CODE SINE(dBm)	97 Hz	997 Hz	10,007 Hz	97 Hz	997 Hz	10,007 Hz	MINOR 2nd	MINOR 6th	OCTAVE	AUGMENTED 11th	MINOR 2nd	MINOR 6th	OCTAVE	AUGMENTED 11th
AKAI S900	40,000	40,010*	WEIGHTED	-98.5	-82.2	-65.9	1.8	.197	.364	.102	.154	.296	.090	.296	.278	.258	.216	.295	.243	.236	.248
			UNWEIGHTED	-89.5	-77.1	-63.3	1.6	.605	.322	.271	.439	.261	.264	.262	.265	.264	.254	.259	.223	.234	.273
AKAI S950	48,000*		WEIGHTED	-101.8	-82.3	-67.9	4.6	.190	.349	.290	.159	.434	.224	.438	.463	.470	.438	.435	.413	.336	.272
			UNWEIGHTED	-92.5	-77.8	-64.0	4.3	.550	.309	.636	.437	.376	.497	.380	.425	.457	.490	.375	.375	.325	.300
AKAI S1000	44,100	44,100*	WEIGHTED	-91.3	-90.1	-85.5	0.3	.072	.143	.053	.054	.091	.052	.126	.085	.077	.098	.125	.102	.098	.099
			UNWEIGHTED	-79.0	-78.9	-77.7	0	.167	.171	.177	.138	.135	.179	.155	.140	.136	.161	.154	.148	.138	.150
CASIO FZ-10M	36,000	35,993*	WEIGHTED	-79.5	-78.5	-73.0	-7.4	.083	.248	1.53	.077	.216	1.50	.281	.299	.337	.458	.263	.386	.186	.280
			UNWEIGHTED	-53.7	-53.7	-53.6	-7.7	.308	.321	2.08	.286	.293	2.03	.345	.387	.410	.594	.353	.424	.293	.452
E-MU EMAX	41,667*	N/A	WEIGHTED	-82.1	-77.1	-63.2	1.5	.268	.992	3.89	.32	1.25	4.01	1.26	N/A	N/A	N/A	1.279	1.28	1.36	1.72
			UNWEIGHTED	-62.7	-62.3	-59.3	1.2	.853	1.012	9.37	1.09	1.20	9.12	1.17	N/A	N/A	N/A	1.195	1.82	2.35	2.96
E-MU EMULATOR THREE	44,100	44,054*	WEIGHTED	-95.8	-88.4	-84.2	6.7	.126	.308	.16	.076	.155	.14	.154	.143	.131	.112	.154	.158	.160	.237
			UNWEIGHTED	-74.3	-72.7	-72.4	6.4	.288	.301	1.36	.147	.185	1.35	.166	.145	.142	.140	.212	.485	.725	1.096
ENSONIQ EPS	48,100	48,077*	WEIGHTED	-84.6	-75.9	-75.5	-6.6	.28	.75	1.16	.20	.47	.40	.47	.46	.44	.40	.47	.45	.43	.38
			UNWEIGHTED	-62.3	-60.0	-59.9	-6.9	4.26	4.28	4.14	4.24	4.24	3.81	4.24	4.25	4.25	4.27	4.24	4.24	4.24	4.24
ENSONIQ MIRAGE MULTI-SAMPLER	29,412	29,412*	WEIGHTED	-78.5	-68.0	-46.7	-1.7	4.8	13.5	27.4	5.1	14.6	27.7	14.34	12.47	11.92	8.70	14.35	14.07	13.77	12.6
			UNWEIGHTED	-63.3	-55.6	-44.8	-2.0	11.8	12.4	35.7	12.8	13.3	35.9	12.90	12.12	12.26	10.32	13.05	12.70	12.96	13.0
KORG DSM-1	48,000	47,690*	WEIGHTED	-87.4	-78.1	-78.2	-4.5	.093	.308	.189	.077	.226	.130	.240	N/A	N/A	N/A	.213	.212	.245	.90
			UNWEIGHTED	-81.2	-75.6	-75.4	-4.8	.235	.267	.367	.174	.195	.243	.209	N/A	N/A	N/A	.188	.188	.337	1.87
KURZWEIL 250	50,000	50,000*	WEIGHTED	-88.3	-78.6	N/A	13.8	1.27	.713	.278	.89	.291	.130	.305	N/A	N/A	N/A	.288	.246	.261	.357
			UNWEIGHTED	-77.7	-74.4	N/A	13.5	3.86	.614	.332	2.89	.253	.191	.266	N/A	N/A	N/A	.251	.222	.244	.384
OBERHEIM DPX-1	N/A	N/A	WEIGHTED	-93.7	-82.3	-72.2	5.4	N/A	N/A	N/A	.78	2.25	.48	2.23	2.00	1.76	N/A	2.25	2.25	2.18	2.00
			UNWEIGHTED	-73.2	-67.4	-65.2	5.1	N/A	N/A	N/A	1.89	1.96	1.19	1.95	1.87	1.75	N/A	1.95	1.97	1.97	1.98
ROLAND S-550	30,000	30,000*	WEIGHTED	-82.4	-82.3	-69.5	7.0	.0624	.094	.0795	.0483	.0614	.0616	.0806	.0777	.0552	.0795	.0864	.0791	.0684	.0791
			UNWEIGHTED	-78.6	-78.5	-68.5	6.7	.1233	.120	.1098	.1023	.1008	.0973	.1157	.1150	.0995	.1140	.1170	.1153	.1080	.1157
SEQUENTIAL PROPHET 2002	41,666	41,677*	WEIGHTED	-98.8	-81.0	-80.7	1.4	.60	1.56	.87	.62	1.88	.85	1.82	N/A	N/A	N/A	1.72	1.64	1.67	1.85
			UNWEIGHTED	-65.1	-64.9	-64.9	1.1	1.80	1.39	4.23	1.84	1.66	4.13	1.60	N/A	N/A	N/A	1.56	1.70	2.05	2.74
SIMMONS SDX	44,100	44,082*	WEIGHTED	-93.0	-90.5	-81.7	1.3	.220	.776	.85	.223	.785	.85	.784	.732	.931	N/A	.779	.765	.736	.708
			UNWEIGHTED	-72.3	-72.2	-71.6	1.0	.679	.672	5.48	.693	.680	5.53	.676	.656	.808	N/A	.682	.767	.844	.959
YAMAHA TX16W	50,000*	49,998	WEIGHTED	-93.0	-93.0	-74.5	-3.8	.077	.159	.103	.0735	.1098	.110	.483	.517	.1109	.696	.229	.204	.1052	.130
			UNWEIGHTED	-76.3	-76.3	-71.9	-4.1	.116	.148	.708	.0991	.1035	.702	.626	.527	.1179	.939	.233	.259	.1046	.153

THD+n (total harmonic distortion plus noise) and dynamic range measurements. ***Sampling rates:*** *Since actual sampling rates are often different from those reported by the manufacturer, where possible we generated our output-stage test sample files not at the reported rate but at a rate derived from our own calculations. (The asterisks in the sampling rate columns show which rate was used to test the output stage of each machine.)* ***Weighted/unweighted:*** *The weighted figures for each instrument have been A-weighted and band-limited at 30kHz. The weighted figures are the ones on which the graphs in this article are based; the unweighted figures are provided for comparison.* ***Dynamic range:*** *For an explanation of the dynamic range figures, see the graph on page 53.* ***Distortion:*** *For an explanation of the distortion figures, see the graph on page 52.* ***Pitch-shift distortion:*** *Distortion was measured while playing back a full-code 997Hz sine wave at various transpositions above and below original pitch. The transpositions of a half-step, a minor sixth, an octave, and an augmented 11th were chosen to present the samplers with a wide variety of mathematical ratios with respect to the original pitch, which different pitch-shift methods will handle differently. The amount of distortion might be higher or lower for other transpositions.*

means that your eyes are blinking open and shut 40,000 times per second. But the world is all darkness; nothing is plugged into your input. Suddenly, you catch a glimpse of a signal coming toward you, and nanoseconds later it's upon you: a vivacious, high-amplitude 10kHz sine wave. You blink one, two, three, four times, barely managing to catch your breath and wipe your silicon brow. "One cycle down," you pant. Four more blinks, and another cycle of the waveform has passed, and four blinks later, another.

Pretty soon, though, you're getting bored. You're in the groove. It's just a little too easy; every four blinks you see the same four voltages over and over again. Zero level, full positive, zero, full negative. Zero, full positive, zero, full negative. Zero, full positive . . . and so on, *ad infinitum*.

"What goes on between zero and full positive, and between zero and full negative?" you start to wonder. "For all I know," you fantasize, "it could be an utterly ideal sine wave, a beautiful curvacious creature of heavenly grace and perfection." A moment later, you realize that this wonderful sine wave may just as well be horribly misshapen; not sinusoidal at all, but ragged and uneven between the four points you're able to see. But you'll never know, because every four blinks the sine wave has completed exactly one cycle. All you can do is dream.

Now consider the case of a sine wave at 10,007Hz. Our heroic analog-to-digital converter, still looking at the fluctuating voltage level 40,000 times per second, will take four samples per period *approximately*, but not exactly. The sine wave will be starting its next period a little before the fifth sample is taken, so the fifth sample won't be identical to the first, the sixth won't be identical to the second, and so on. With each successive period, the converter will capture a new set of four voltages, with the result that over a large number of periods, it will see most of what goes on in the sine wave—peaks, zero crossings, distortions, everything.

By skewing the conventional test frequencies of 100Hz, 1kHz, and 10kHz, we managed to dodge the possibility that one sampler or another might be blind to distortions ocurring between full positive and full negative amplitude. Just to be sure, we divided 97, 997, and 10,007 into the sampling rate of each instrument, assuming that if they didn't divide evenly, we were in the clear. Much to our relief, these relatively prime numbers worked in every case. In fact, using these values killed two tweeters with one voltage spike: It allowed us to avoid biasing our tests in favor of one sampling rate or another, and to use the same frequencies for testing all sampling rates.

ALL THE WORLD'S A STAGE

AT THIS POINT, WE'RE ALMOST READY to take a spin through the wild and wacky world of sampling rates. But before we do, it would be helpful to firm up the distinction between a sampler's input and output stages. Simply put, the recording part of a sampler is its input stage, and the playback part is its output stage.

Why should anyone care? Consider this. It's generally acknowledged that sampling your own sounds is a lot of trouble, particularly if you want conventional instrument sounds such as a string section or a piano. Many musicians use this sort of sound exclusively, and, of course, they don't record the samples themselves. They buy them from a developer. As a result, they may never use the input section of their instrument. In fact, neither may the developer. Many developers take advantage of the superior input performance of a high-end sampler, porting the sample data over to lower-grade machines via MIDI. If you rely exclusively on third-party or factory samples, you may not care so much about the input-stage performance of your machine.

Conversely, if you're a developer, or if you're prone to recording, editing, and keyboard-mapping your own sounds, and perhaps using them in several samplers, the input stage of your main axe may be much more critical than the output stage. Of course, musicians who tend to use the same instrument for both sampling and playback are most interested in how well signals fare after passing through both stages.

The MIDI sample dump and system-exclusive specifications provide an ideal way of isolating a sampler's input and output processes for testing. Via MIDI, samples can be ported out of an instrument after having been digitized by the input stage, but before conversion back into audio via the output stage. Likewise, sample files created entirely in software—without having been subjected

to the distortions of an analog-to-digital conversion—can be ported into the sampler and played back through the output stage.

Digidesign's software product line, which includes Sound Designer Universal, Softsynth, and Turbosynth, was critical in establishing a set of standardized tests for measuring output-stage and input-plus-output-stage performance. Universal Sound Designer served as our basic environment for looping, viewing, and porting sampled material. The resolution of Sound Designer's sample file format is 16-bit. Unfortunately, the program truncates to 12 bits for data dumping. Blank Software's Alchemy program, which performs many of the same functions as Sound Designer, maintains 16 bits while performing MIDI dumps, so we ended up using Alchemy for porting samples to and from 16-bit instruments. (As it happens, the Akai S950 understands the S900's sys-ex dump format as implemented by Alchemy, but not as implemented by Sound Designer. Therefore, we used Alchemy with the 12-bit S950 as well.)

Digidesign's Softsynth was used to compute 16-bit, Sound-Designer-compatible sine wave sample files at 100% amplitude. (This is the highest level short of digital clipping, and is known as *full-code*.) Such files had to be generated at the selected test frequencies and at the proper sampling rate for each machine. We also used Softsynth to generate white noise for our output-stage frequency response tests. Turbosynth also generates sine waves and noise, but programmer Mark Jeffery, who invented both programs, suggested that the sine and noise algorithms in Softsynth are closer to ideal. Turbosynth's square waves, however, are squarer than Softsynth's, so we used Turbosynth for our output-stage transient response tests. Armed with a set of test signals which had never seen the horror of analog-to-digital conversion, we were prepared to measure the performance of the samplers' output stages in isolation.

THE TROUBLE WITH TABLES

BOB CHIDLAW, CHIEF SYSTEMS ENGINEER at Kurzweil, found reason to question the quality of the Softsynth sample files. He points to evidence in the K250's output-stage distortion plots of "some low-level, non-harmonic distortion" which doesn't appear to have come from the instrument itself. The distortion, he continues, is "probably due to using a table look-up method" of computing digital sine waves, which is in fact the method employed by Softsynth. "To do better would take significantly more computation time," Chidlaw continues. "Digidesign probably made the right compromise here for what they wanted their product to do, but it is not really ideal for use in a laboratory situation." Mark Jeffery concurs with this observation.

The table-lookup distortion varies depending on the sine wave's frequency and sampling rate. Given that Softsynth is our best source of output-stage test files, such distortion is an unavoidable result of our decision to use the same frequencies to test every instrument. Distortion originating in the Softsynth files is often recognizable by its characteristic double peaks, which are prominent in the plot below.

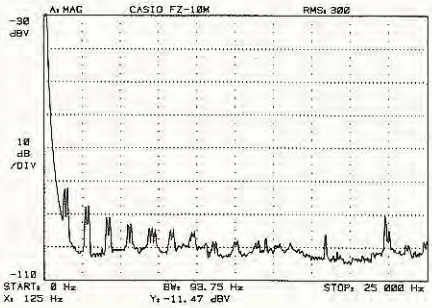

However, double-peaked distortion does show up occasionally in input-plus-output plots as well, so it's hard to be sure of its origin.

Evaluating the performance of the input stages alone proved to be more problematic. The H-P 3561A, the Tektronix AA 501, and the other test gear we used all require an audio input signal. Post-input-stage signals, having been digitized and stored in memory, exist not as audio but as binary code. Turning them into audio would involve passing them through a digital-to-analog converter, which would introduce its own distortions. How could we analyze the effect on our test signals of the input stages only?

Rick Downs of Burr-Brown, Inc., a manufacturer of A/D converters and DACs,

AKAI S900

Word has it that more musicians own S900s than any other sampler on the market. If specs are any indication, it's not hard to tell why. Depending on how you want to look at it, the Akai S900 ranks either just behind the best samplers, or just ahead of the instruments on the next rung down.

Frequency Response. As we've noted elsewhere, the S900's input-plus-output-stage frequency response is far from flat. The low end drops by several dB below 500Hz (this is more discernible in the logarithmic-scale plot, not shown here). More prominent is the boost of nearly 5dB between 3.5kHz and 16kHz, beyond which it rolls off, falling 6dB below the level at 1kHz by around 17.5kHz. We'd bet that this high-frequency bump has something to do with the S900's popularity, even though it makes the instrument less useful for applications requiring high fidelity. The output-stage frequency response is very nearly flat until it begins to roll off at around 14kHz.

Dynamic Range. When it isn't playing anything, the S900 is among the quieter samplers, and its output-stage residual noise is about average. But with both the input and the output stages grinding away—the most important measurement, in our view—the S900's noise floor comes close to those produced by the two noisiest instruments, the E-mu Emax and the Ensoniq Mirage. Incidentally, the silent sample used in the input-plus-output residual noise test, like all of our S900 samples, was recorded with the input level pot at around 10 o'clock. When the input pot is more than halfway up, the internal noise level is so high that it actually registers on the instrument's LCD meter!

Distortion. Averaged, the S900's weighted THD+n figures place it just behind the E-mu EIII, along with the Korg DSM-1 and the Akai S950. In the 10,007Hz tests, the S900 actually did better than the EIII. Distortion components that look suspiciously like products of Softsynth's computational method appear in the output-stage plots, which could mean that output-stage distortion is even better than our measurements indicate.

Pitch-Shift Distortion. As one would expect of a good variable-rate playback system, the S900's unweighted pitch-shift distortion measurements change little regardless of the interval played. The plots show that the distortion components already present in the 997Hz sample simply spread out when shifted upward, and compress when shifted downward. In plots depicting downward transpositions, the effect of a tracking filter can be seen. It keeps imaging components out of earshot.

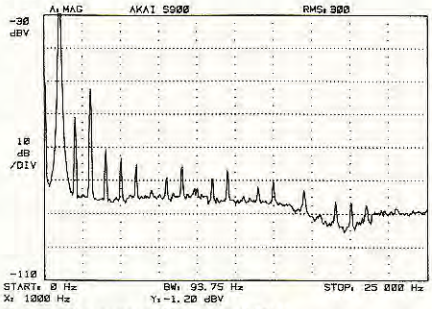

Input-plus-output distortion, 997Hz.

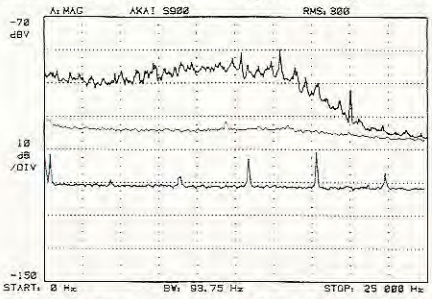

Residual noise: Input-plus-output (black), output-stage (red), quiescent (blue).

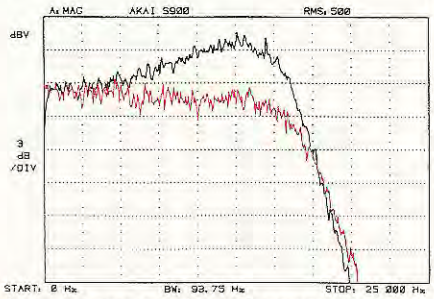

Frequency response: Input-plus-output (black), output-stage (red).

AKAI S950

You gain a little here, you lose a little there. The Akai S950, an update of the S900, differs from its predecessor in offering sampling rates up to 48kHz. The primary benefit is a bandwidth of 21kHz—nearly the broadest found. On the other hand, the S950 has somewhat higher distortion levels than the S900. It also produces a strange clicking when playing back high frequencies, as in extreme upward transpositions and our 10,007Hz tests. (We found a similar quirk in the Simmons SDX.)

Frequency Response. Like the S900, the S950 introduces a high-frequency emphasis at the input stage. This boosts frequencies between 7kHz and 21kHz by 5dB. The output stage is nearly flat to 17kHz with a not-too-leisurely cutoff, making it ideal for playback of samples imported from other machines.

Dynamic Range. The S950's output is 2.8dB hotter than that of the S900, but otherwise it exhibits the same dynamic range characteristics. Quiescent noise and output-stage residual noise are relatively low, and input-plus-ouput residual noise is a little high for comfort.

Distortion. The S950's distortion rating falls short of the high standard set by the S900, though not by much. (The plots suggest that the increased output-stage distortion at 97Hz and 997Hz may have more to do with Soft-synth than with the sampler, but we can't be sure.) The 10,007Hz plots show remarkably few distortion components, but illustrate the S950's difficulty in handling high-frequency material: The basic noise floor is unusually flat, more than 10dB higher than that produced by the S900 in the same test, and covered with a fringe of tiny ridges. To the ear, this translates into something halfway between buzzing and clicking.

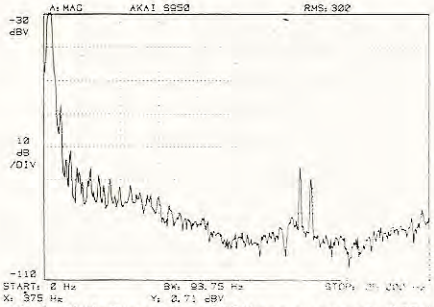

Pitch-Shift Distortion: 997Hz an augmented 11th downward.

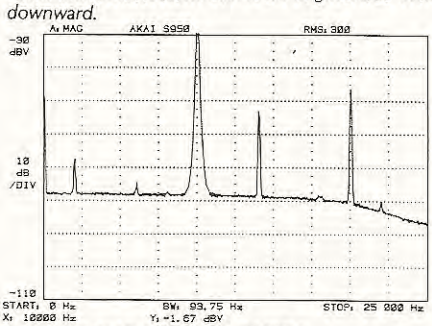

Input-plus-output distortion, 10,007Hz.

Pitch-Shift Distortion. Generally, the S950 yields pitch-shift plots identical to those of the S900. Distortion peaks spread out elegantly during upward transposition and compress when pitch is shifted downward. However, the octave-upward plot shows the same fringed noise floor as the 10,007Hz tests. The effect is more pronounced in the plot of the augmented-11th transposition.

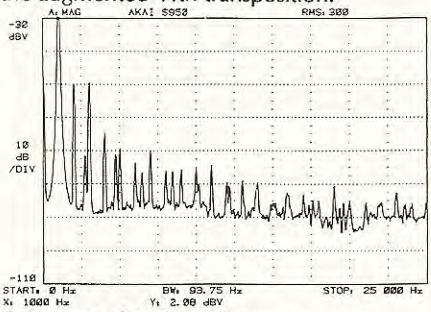

Output-stage distortion, 997Hz.

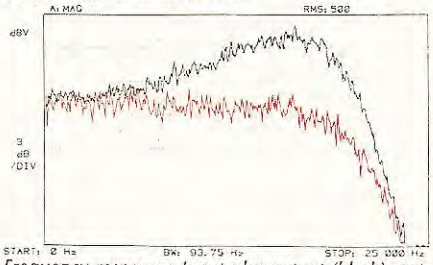

Frequency response: Input-plus-output (black), output-stage (red).

AKAI S1000

How about that Akai S1000? As confirmed in our listening tests, the S1000 is the prime contender for overall best. It combines the quietest silent-sample playback, the lowest overall distortion average, one of the top pitch-shift distortion ratings, and one of the broadest bandwidths. All specs considered, the S1000's only real competition comes from the E-mu EIII and—here's a monkey wrench for you—the Korg DSM-1.

Frequency Response. There are flatter frequency response curves out there, but the S1000's looks pretty good. The combination of input and output stages exhibits an ever-so-slight bump between 5kHz and 16kHz; the output stage alone rolls off very gradually above that range until the cutoff, at a little over 22kHz. The cutoff slope is steep enough that aliasing isn't an issue.

Dynamic Range. The S1000's output level and quiescent noise floor are average. Its output-stage residual noise level ranks third (behind the Yamaha TX16W and the Simmons SDX). The residual noise of the S1000's input and output stages together edges out that of the EIII by 1dB.

Distortion. The only instrument for which we measured a lower distortion average is the Roland S-550. (The Yamaha TX16W ran a close third, after which the margins become a bit wider.) In the input-plus-output-stage tests, the THD+n figures we're talking about are enticingly close to those of the original test signals, and the spectrum plots show it. As you might expect, the output stage alone is even cleaner. Both the output stage alone and the input-plus-output combination measure less distortion at 997Hz than at either 97Hz or 10,007Hz, which seems odd to us, but is consistent with results obtained from other high-performance machines.

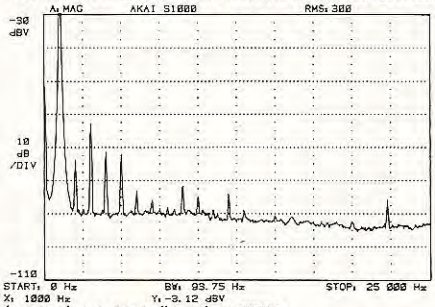

Input-plus-output distortion, 997Hz.

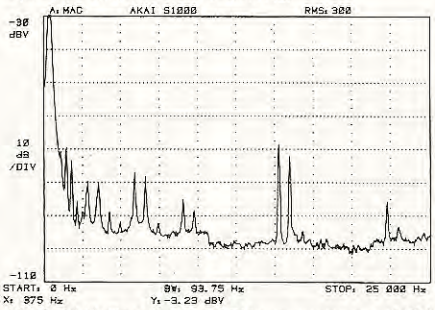

Pitch-Shift Distortion: 997Hz an augmented 11th downward.

Pitch-Shift Distortion. The S1000's variable-rate clock keeps pitch-shift distortion to a minimum. The unweighted THD+n measurements show a negligible degree of variation regardless of transposition, and the plots confirm that the distortion components change slightly in amplitude, and not at all in frequency, as pitch is shifted.

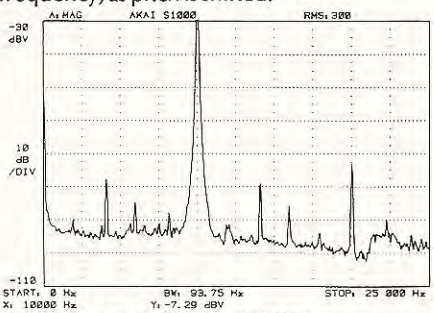

Input-plus-output distortion, 10,007Hz.

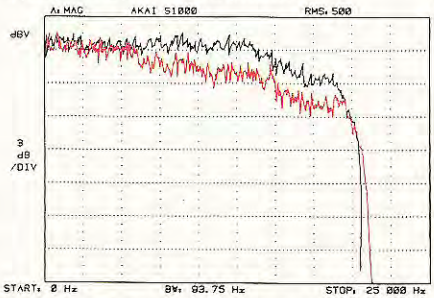

Frequency response: Input-plus-output (black), output-stage (red).

CASIO FZ-10M

After two years on the market, the FZ-10M remains the lowest-priced 16-bit sampler. The extra resolution, however, doesn't make it a clear winner. A high THD+n at 10,007Hz knocks it out of the low-distortion category into which it would otherwise fall. Its quiescent noise floor and output-stage residual noise rank near the bottom of the heap. As for frequency response, it's extraordinarily flat, but drops off at 15kHz—just a little early, given a sampling rate close to 36kHz.

Frequency Response. We wish all samplers had frequency responses as flat as the FZ-10M's input-plus-output curve. The anti-aliasing filter's cutoff is as close to brick-wall as we've seen. The output-stage curve is nearly as good.

Dynamic Range. The FZ-10M does poorly in the quiescent and output-stage residual noise tests, but the combination of input-and output-stage noise lands it smack-dab in the middle. However, the FZ-10M's low output level—the lowest we encountered—is bound to result in a higher system noise floor when the instrument is used in conjunction with a mixer.

Distortion. If we were to exclude the 10,007Hz tests, the FZ-10M's distortion average would compare well with that of the E-mu EIII and Korg DSM-1. The spectrum plots, which are remarkably clean at 97Hz and 997Hz, bear this out. Including 10,007Hz, at which THD+n measures a hefty 1.5% (weighted), the instrument's standing drops to the general level of the Simmons SDX, Kurzweil K250, and Ensoniq EPS. Indeed, the plots show a number of high-amplitude, mostly inharmonic components on either side of 10,007Hz. If you tend to play orchestral sounds, the 10kHz range matters much less than if you sample rock-and-roll CDs.

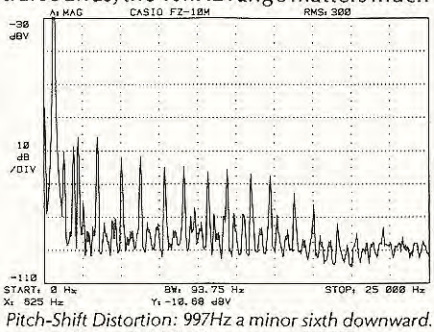

Pitch-Shift Distortion: 997Hz a minor sixth downward.

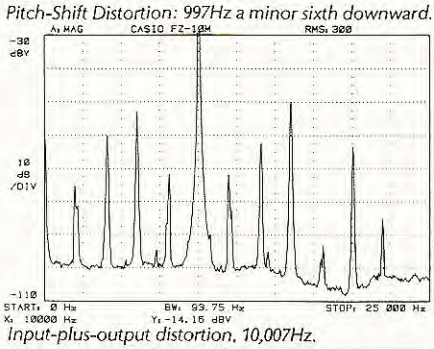

Input-plus-output distortion, 10,007Hz.

Pitch-Shift Distortion. The FZ-10M's pitch-shift plots look disturbingly different at various transpositions. A half-step upward shift brings out inharmonic distortion; an octave up, the distortion is exclusively harmonic. Although Casio reports a variable-rate clock, these are characteristics of a fixed-rate system. Imaging isn't a problem.

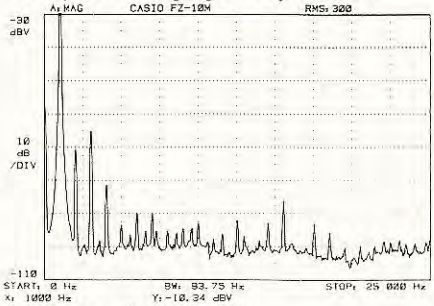

Output-stage distortion, 997Hz.

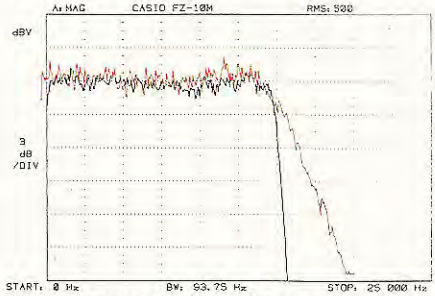

Frequency response: Input-plus-output (black), output-stage (red).

E-MU EMAX

The Emax, which has been well-received by musicians, isn't among the better performers in our survey. In nearly every area, it ranks above only the Ensoniq Mirage. Bandwidth, which spans 18kHz, is one bright spot. Another is the Emax's ability—particularly that of the output stage—to reproduce a square wave.

Frequency Response. The Emax's frequency response curves are much like those of the EIII and the Simmons SDX: a gentle slope extending beyond the Nyquist limit for the output stage, with a steep cutoff below Nyquist thrown in for the input-plus-output response. Also as with the EIII and SDX, imaging appears above the anti-aliasing filter's cutoff in the input-plus-output plot.

Dynamic Range. In the playback of both zero-amplitude words (output-stage residual noise) and sampled silence (input-plus-output residual noise), the Emax beats only the Mirage.

Distortion. The doubled peaks we've come to expect from Softsynth appear at 97Hz in the input-plus-output plot as well as the output-stage plot, leading us to suspect that, in this case at least, they originate with the sampler. Harmonic peaks buttressed by inharmonic components extend across the frequency range in the 997Hz plots, and likewise with inharmonic peaks in the 10,007Hz plots. (E-mu's Dave Rossum attributes much of this distortion to the Emax's data compression scheme.) The Emax's input-plus-output distortion is lower than that of the output stage alone, which is unusual. We'd venture a guess that this is because our input-plus-output samples aren't as close to full code as ideally they ought to be.

Pitch-Shift Distortion. The Emax's fixed-rate algorithm produces mostly harmonic, but inconsistently distributed, distortion.

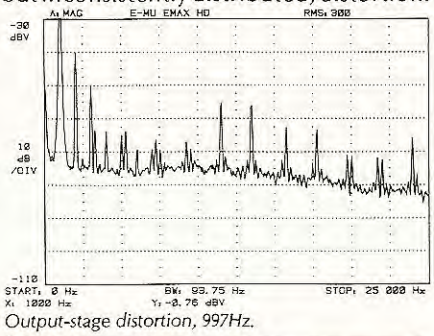

Output-stage distortion, 997Hz.

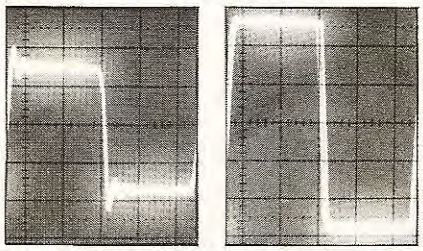
Transient response: Input-plus-output, output-stage.

(Upward transpositions beyond a half-step aren't available at the highest sampling rate.) For a fixed-rate system, the imaging revealed by downward shifts is remarkably like that produced by a variable-rate scheme. According to Rossum, the greater the oversampling, the more the mathematics of fixed-rate transposition approach those of variable-rate transposition.

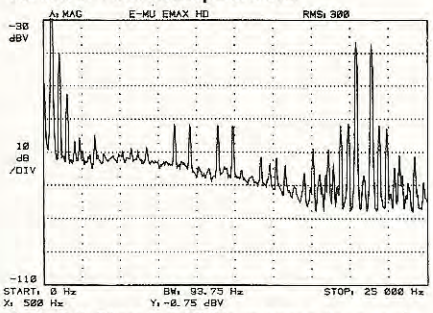

Pitch-Shift Distortion: 997Hz an octave downward.

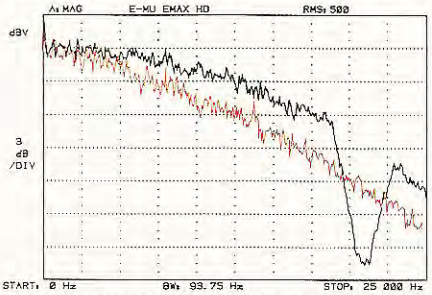

Frequency response: Input-plus-output (black), output-stage (red).

offered a solution. Rick wrote a software routine that translates Sound Designer's file format into a format readable by DaDisp, a scientific signal analysis package for the IBM PC (manufactured by DSP Development Corp). DaDisp is capable of analyzing the spectra of digitally encoded sound *within the digital domain*. For example, here's a plot of our 997Hz tone directly after passing through the input stage of the Yamaha TX16W:

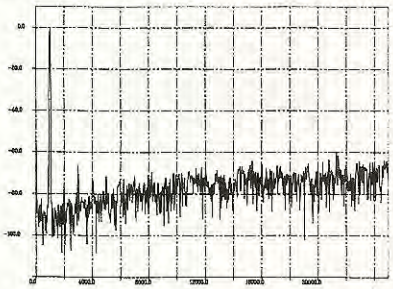

Unfortunately, the task of generating some 160 plots (10 per instrument) was too big to be completed in time. Input-stage results for several instruments weren't available prior to press time, and some of the differences between DaDisp's plots and those of the H-P 3561A analyzer suggest that their computational algorithms may have been set up too differently for our purposes. (For instance, we used a Hanning window with the H-P 3561A, while Rick set up his equipment with a Blackman-Harriss Minimum-Term Fourth-Order window—but that's another story.) Rick was able to produce direct-from-digital plots of some of our output-stage test files, which were helpful in evaluating Bob Chidlaw's concern regarding distortion in the Softsynth sine waves. In the end, however, not enough of Rick's results were available to allow us to draw certain conclusions about their meaning. Therefore, only the output-stage and input-plus-output results are considered in this article.

Q: WHY DID THE SAMPLING RATE CROSS THE ROAD?

ONCE WE HAD DECIDED TO TEST INput and output stages separately, it was time to generate a batch of Softsynth test files for each instrument's sampling rate. But which sampling rate? Many instruments offer more than one. On the assumption that musicians would rather have wider bandwidth than extra sample memory, we decided to test each sampler at the highest available sampling rate. Little did we realize that life in the digital domain isn't so simple.

Assuming you want to find out a sampler's highest sampling rate, where do you look? In the user's manual, of course.

But there's a catch: When you feed an instrument a sample file cooked up in Softsynth at 997Hz according to the sampling rate reported in the instrument's manual, as often as not the instrument will spit out some frequency other than 997Hz. (This tidbit thanks to our trusty Hewlett-Packard 5316B Universal Counter.) Which means one of two things: Either Softsynth is out to lunch, or the sampling rate isn't what the manual says it is. Years of painful experience indicate that the latter is the correct assumption.

The search continues. Some machines display their sampling rates in their front-panel LCDs. As it turns out, the figure found there is often different from the one printed in the manual. When we tried putting the LCD's number to the test, it, too, turned out to be wrong as often as it was right.

Becoming incredibly crafty, we realized that a sampling rate value is encoded into the header of Sound Designer files, and that the "get info" option in that program would allow us to see it. We recorded a sample and ported it over to the Macintosh, eagerly choosing "get info" from the menu. Wrong again. The sample-rate header is encoded by the manufacturer, so it's no more likely to be correct than the manual or the front-panel display.

Next option: Consult the manufacturer. We're happy to report that this technique works beautifully with American and British companies. Japanese companies, however, are another story. Either nobody knows the exact sampling rates, or nobody's telling, or they simply don't understand the question. Backed into such a tight corner, we were left with only one course of action—investigative journalism.

We attacked each machine in turn, applying a handy algebraic formula as follows (listen closely): the frequency entered into Softsynth divided by the sampling rate entered into Softsynth equals the frequency

E-MU EMULATOR THREE

As the most expensive instrument on our roster, the E-mu Emulator Three is accompanied by great expectations. Although it doesn't fulfill them entirely, it does leave the impression that the extra dollars do, indeed, buy quality. In overall performance, the EIII ranks second to the Akai S1000. Its frequency response is flatter, its average distortion a bit higher, its pitch-shift scheme slightly less effective, its residual noise levels a tad higher, and its output a good deal hotter.

Frequency Response. The EIII's input-plus-output frequency response curve is probably the most accurate we saw. Flat up to 15kHz, it rolls off gradually before dropping steeply a little beyond 20kHz. The output-stage curve looks just as good, but continues its slope without a sharp cutoff. According to E-mu's Dave Rossum, the range between 23kHz and 25kHz that appears to have been left behind by the anti-aliasing filter actually represents imaged white noise left behind by the reconstruction filter. This isn't a problem, judging from the pitch-shift plots.

Dynamic Range. The EIII's hot output level is complemented by very low noise floors. Quiescent and output-stage residual noise levels are well above average, and the silent sample plays back a good deal more quietly than that of every other instrument except the S1000.

Distortion. Three instruments beat the EIII's distortion average: the Roland S550, the S1000, and the Yamaha TX16W. This puts it on equal footing with the Korg DSM-1, and only slightly ahead of the Akai S900. The EIII's output-stage THD+n figures are better than those for the combined input and output stages, even though the output-stage plots look messier. This suggests that Softsynth might have played a role in the output-stage distortion figures, as low as they are. (The general noise floor in the input-plus-output plots is clearly higher, which confirms the numbers.)

Pitch-Shift Distortion. When it comes to the *average amount* of distortion added by transposition (weighted), the EIII looks better than everything except the Kurzweil 250. The unweighted THD+n measurements, however, vary by a good deal, rising steadily as pitch is transposed downward.

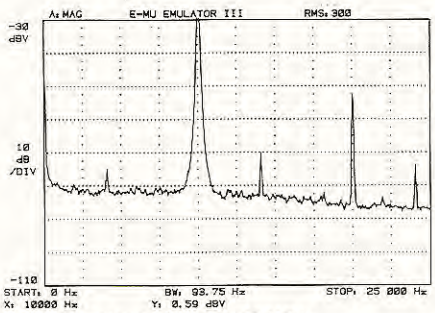

Input-plus-output distortion, 10,007Hz.

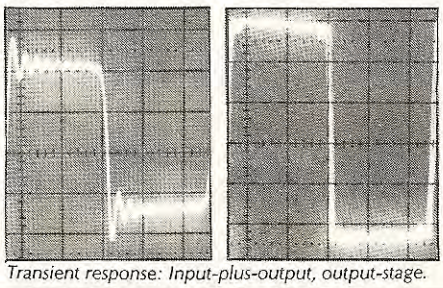

Transient response: Input-plus-output, output-stage.

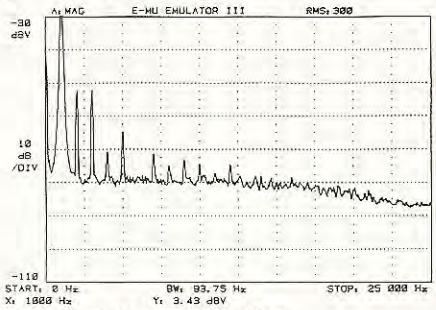

Input-plus-output distortion, 997Hz.

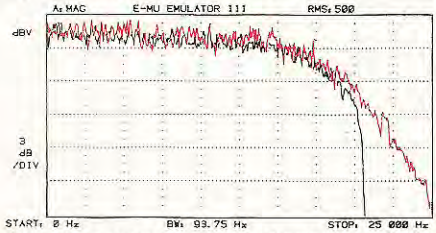

Frequency response: Input-plus-output (black), output-stage (red).

generated by the sampler when playing back the Softsynth file divided by the true sampling rate. It worked. Aside from the sampling rates in question, this is what we learned: Sampling rates are rarely nice round values such as 36kHz or 44.1kHz. Often they have repeating decimal values that extend to infinity. Furthermore, it appears that variations in clock speed from instrument to instrument of the same make and model aren't unheard of. We tested two different Akai S900s, one fitted with the Marion Systems 16-bit upgrade and one in its off-the-shelf state. In contrast to the reported sampling rate of 40kHz, the Marion-equipped S900 clocked in at around 39,996Hz, and the vanilla model at around 40,010Hz. (Just in case you're wondering, these are very small discrepancies, more a curiousity than anything worth worrying about.)

When making up our output-stage test files in Softsynth, then, we used the experimentally derived sampling rate rather than the ones to be gotten from the manual, the LCD, the sample-dump header, or the horse's mouth.

A: TO KEEP ITS PLAYBACK RATE UP!

THERE ARE THREE EXCEPTIONS, HOWever. The Ensoniq Mirage and EPS have *sample playback rates* that are lower than their highest sampling rates. Digital audio theory holds that a sampler will perform best when sampling rate and sample playback rate are equal, so we tested the Ensoniq instruments at the lower sampling rates, experimentally derived.

The third exception is the E-mu Emax. Its playback clock runs a great deal faster than the sampling clock, allowing it to interpolate new sample values during pitch-shifted playback rather than doubling or dropping samples. (This produces lower pitch-shift distortion than a simple drop/add-sample scheme.) However, the Emax's playback rate bears no rational mathematical relationship to either of its two sampling rates. This makes its sampling rate impossible to determine algebraically. Therefore, we accepted the reported sampling rate of 41,666.66Hz (rounded to 41,667Hz for the sake of Softsynth, which doesn't accept noninteger sample rate values).

IMAGING OVER SUBSTANCE

BOTH THE H-P 3561A, WHICH GAVE US our spectrum plots, and the Tektronix AA 501, which measured THD+n, work by averaging their measurements over several seconds' time. In order to allow them to do their jobs, it was necessary to loop the test samples. Looping itself, however, brought a new family of digital gremlins out of the woodwork.

We've already mentioned aliasing and imaging, the two classes of distortion inevitably produced by the mathematics of sampling. Aliasing, being an input-stage phenomenon, doesn't affect samples generated in Softsynth. Imaging, on the other hand, is quite evident in our Softsynth files. In the analog domain, imaging appears as supersonic components of the kind seen in the output of the Ensoniq EPS. In the digital domain, it usually takes the form of subaudio amplitude modulation. In fact, such modulation can be seen in the waveform displays of Sound Designer, Alchemy, and other sample editing programs if you zoom out far enough. Consider the 10,007Hz-sampled-at-40kHz situation described earlier:

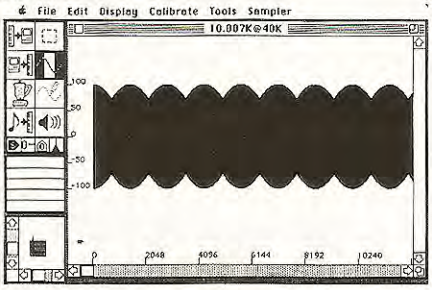

As you can see, such modulation patterns are periodic, usually on the order of several cycles per second, sometimes slower. Close up, it's easy to see how the changing shape of the waveform produces this kind of pattern.

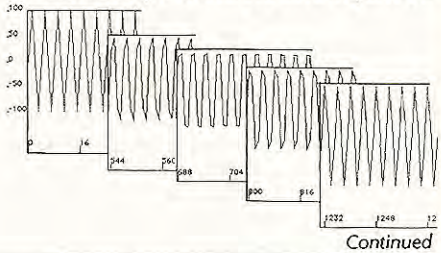

Continued

ENSONIQ EPS

The EPS is an interesting case. Its true THD+n figures are outrageously high, due to distortion above the audio range. Adjusted to better reflect what the instrument actually sounds like, the figures are quite low. Another funny thing: The EPS performs better at its 48,077Hz sampling rate than at 52,082kHz *when the number of oscillators is set to 12*. That's because the 12-oscillator setting also sets the playback rate to 48,077Hz, and in digital audio, parity between sampling and playback rates produces optimal performance.

Frequency Response. The EPS scores better than average in this area. The 48kHz sampling rate allows a bandwidth of nearly 24kHz—greater than that of the E-mu EIII and Akai S1000—although in practice the EPS rolls off at about 19kHz. Both input-plus-output and output-stage frequency responses are broad, relatively flat, and roll off gently. However, the supersonic output components revealed in our tests indicate that the reconstruction filter is a bit lax. Even though these frequencies are beyond the range of human hearing, their interactions with other audio equipment could, under some circumstances, result in perceptible distortion.

Dynamic Range. Given that the EPS's output level is nearly the lowest we encountered (-6.6dBm weighted), the noise floor measurements must be considered a bit worse than they look. Unfortunately, they don't look that great to start with. The EPS is at the bottom of the list for quiescent noise floor and output-stage residual noise. It jumps into the middle of the pack, surrounded by the TX16W and the DSM-1, in the playback of sampled silence. Sampled silence plays back no more noisily than zero-amplitude words, which is one of the hallmarks of all-digital designs.

Distortion. Generally speaking, the EPS has more trouble with high-frequency inputs than low. The input-plus-output THD+n figures show that distortion increases with the frequency of the test tone. At 10,007Hz, distortion—much of it inharmonic—is over 1% (weighted), a figure exceeded only by the CasioFZ-10M. The EPS's output stage is nearly twice as distortion-free, on average, as its input stage. This suggests that the instrument is much better for playing back third-party samples than for creating your own.

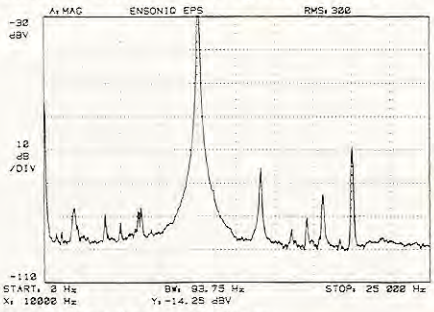

Output-stage distortion, 10,007Hz.

Pitch-Shift Distortion. The EPS, which employs fixed-rate playback, performs mathematical interpolations in order to shift pitch. The effectiveness of Ensoniq's algorithm is demonstrated by the extreme consistency of pitch-shift distortion throughout the intervals tested. The plots of downward shifts are free of the transposed imaging characteristic of variable-rate schemes, and the upward-shift plots look equally good.

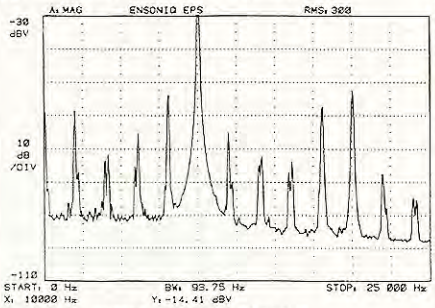

Input-plus-output distortion, 10,007Hz.

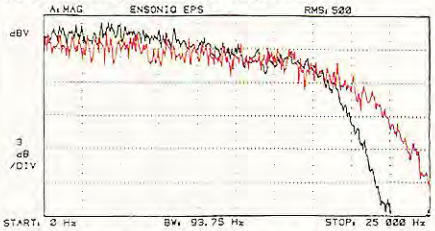

Frequency response: Input-plus-output (black), output-stage (red).

KORG DSM-1

According to the grapevine, the Korg DSM-1 can be found brand-new for under $1,500. Our tests indicate that that just might be a deal. The only area in which the DSM-1 is clearly "best" is bandwidth, which spans 22.5kHz. It turns in a solid showing in just about every other test category, though, putting it head to head with the Akai S950 as a contender for third place (behind the Akai S1000 and the E-mu EIII).

Frequency Response. Okay, so it reproduces frequencies you won't be able to hear by age 40. But is it flat? About like a roller-coaster ride. The output stage's frequency response drops around 5dB between 0Hz and 7.5kHz, flattens out for a while, makes a 3dB ski-jump between 18kHz and 22kHz, and then plunges sharply. The input-plus-output curve begins with a similar drop, but then rises 6dB between 7.5kHz and 18kHz. From there, it takes the 3dB ski-jump before cutting off for real at 22.5kHz. What does all of this mean? Certainly, output-stage response strengthens the lows at the expense of the highs. As for the input-plus-output response, it might add brightness and punch, but might also weaken the critical range between 3kHz and 8kHz.

Dynamic Range. The DSM-1's output level is low, which doesn't particularly complement its high quiescent noise floor. In operation, where it really counts, it looks much better. The fact that the input and output stages together are only a few dB noisier than the output stage alone is a good sign.

Distortion. When averaged, the DSM-1's weighted THD+n figures are neck-and-neck with those of the EIII, and just behind those of the Roland S-550, the Akai S1000, and the Yamaha TX16W. At 997Hz, though, the DSM-1's THD+n is double that of the other instruments. The input-plus-output plots look nearly identical to those of the output stage alone: very tidy. The instrument's low output level, however, makes the actual noise floors higher than they look.

Pitch-Shift Distortion. The DSM-1's variable-rate clock maintains reasonably low and consistent distortion over most of the intervals we tested. (At the 47,690Hz sampling rate, upward transposition is limited to a perfect fourth.) Imaging components made it past the output filter only in the augmented-11th downward shift, showing up an unfortunate 36dB below the fundamental.

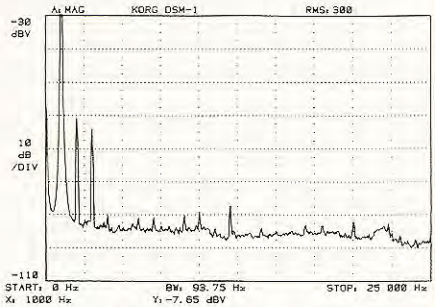

Output-stage distortion, 997Hz.

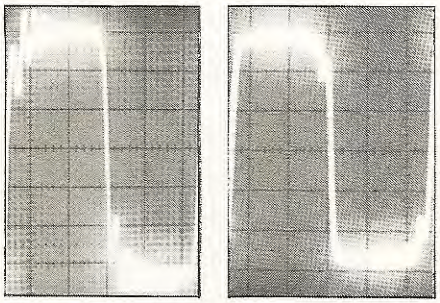

Transient response: Input-plus-output, output-stage.

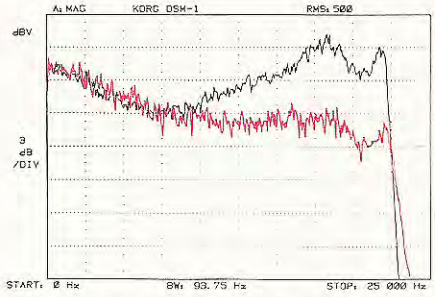

Frequency response: Input-plus-output (black), output-stage (red).

ENSONIQ MIRAGE MULTI-SAMPLER

Dog meat, right? Well, it's hard to deny that the Mirage's results look pretty horrendous. However, your humble reporter must interject that he likes this machine a lot. The grungy sound is part of its charm, and the MASOS user interface, while cryptic, has an attractive flexibility and logic to it. Like the old VW Beetle, the Mirage is ugly, but it gets you there. Are we making up excuses for a sampler that failed our tests miserably? You bet. On the used market, you might be able to pick up five of these babies for the price of a new 12-bit machine, and given the heroic efforts of third-party developers to provide good-sounding samples for the Mirage (K-Muse *et al.*), you might be better off this way.

Frequency Response. In both the output and input-plus-output plots, the Mirage's frequency response drops steadily from around 2kHz on. The point at which the output-stage curve falls below the level at 1kHz by 6dB—a generous reckoning of bandwidth—is at approximately 8.5kHz, surprisingly early given a 29,412kHz sampling rate. The -6dB point for the input and output stages together is around 12kHz.

Dynamic Range. The Mirage is at the bottom of the list in all three noise floor measurements.

Distortion. The Mirage's distortion levels surpass those of any other sampler by such a wide margin that it may as well be considered a synthesizer. One interesting feature is the degree of pitch inaccuracy in the 10,007Hz tests. The H-P 3562A usually quantizes 10,007Hz to 10kHz; the Mirage's plots show 10,187Hz. Ensoniq's Bill Mauchley comments, "The old hi-fi axiom that specs don't tell the whole story is especially true in musical instruments. The Mirage, though it does a bad job of reproducing 10kHz sine waves, is certainly a viable instrument."

Pitch-Shift Distortion. Thoroughly inconsistent, with gobs of inharmonic distortion throughout the range. As with the 10,007Hz tests, upward shifts beyond an octave exhibit a degree of pitch inaccuracy.

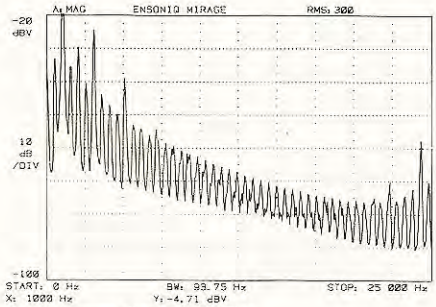

Output-stage distortion, 997Hz.

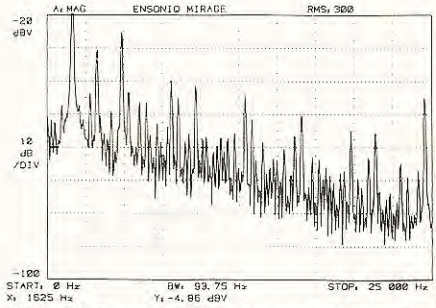

Pitch-Shift Distortion: 997Hz a minor sixth upward.

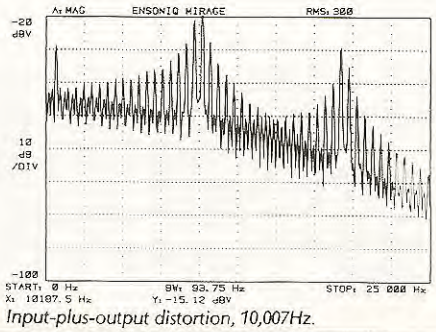

Input-plus-output distortion, 10,007Hz.

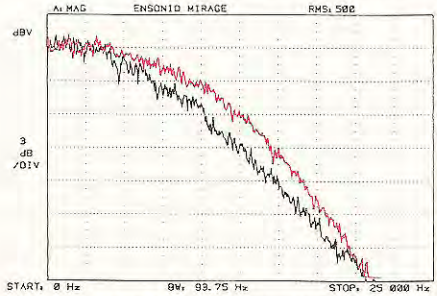

Frequency response: Input-plus-output (black), output-stage (red).

KURZWEIL 250 w/SAMPLING OPTION

Rather than simply recording a signal, the K250's "sound modeling" technique normalizes its amplitude variations. Then it applies a 256-stage amplitude envelope whose contour is based on that of the original sound. This keeps all 10 bits of resolution active while the sampler is playing, eliminating the quantization noise that would otherwise occur as the sound dies away. (It also makes the instrument more practical for playing back factory samples than for rolling your own.) It's possible to bypass these manipulations—and the improvements in signal-to-noise ratio—by using what Kurzweil calls de-emphasis mode. We tested the K250 in this mode.

Frequency Response. Why, given a 50kHz sampling rate, do the K250's filters roll off at 16kHz? We assume the answer is to accommodate lower sampling rates, but without adjustable cutoffs of the kind found on the Ensoniq EPS, a good 8kHz of bandwidth goes down the drain. The output-stage curve falls gently from 2.5kHz on; the input-plus-output response features either a bump between 2.5kHz and 12kHz, or a dip between 500Hz and 6kHz, depending on how you look at it. Both drop off below 100Hz or so (more apparent in a logarithmic plot.)

Dynamic Range. The K250 is hard to evaluate in this regard, since the sound modelling process renders our input-plus-output residual noise measurement meaningless. Its output level is the hottest we saw.

Distortion. While the K250's average distortion comes very close to that of the EPS, the K250 differs in that its 97Hz tests measure the most distortion and its 10,007Hz tests the least. The output stage is cleaner than the input-plus-output stages by nearly a factor of two but, oddly, the difference isn't discernible in the plots. By and large, the 97Hz plots have higher-than-average noise floors. The 10,007Hz plots have next to no distortion, but their noise floors rise surprisingly high in the low range.

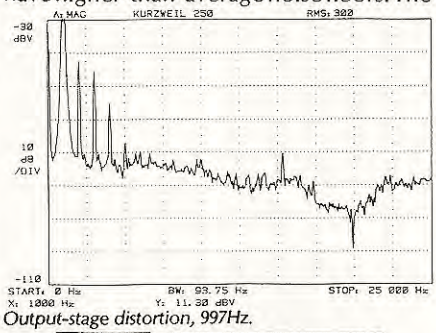

Output-stage distortion, 997Hz.

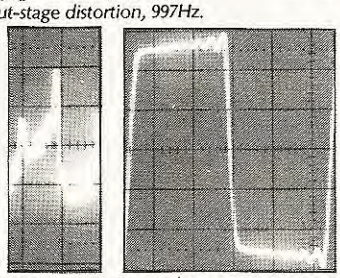
Transient response: Input-plus-output, output-stage.

Pitch-Shift Distortion. In addition to the extraordinary effectiveness of the K250's variable-rate playback scheme, the pitch-shift plots reveal the effect of a strong tracking filter. The filter eliminates all transposed images.

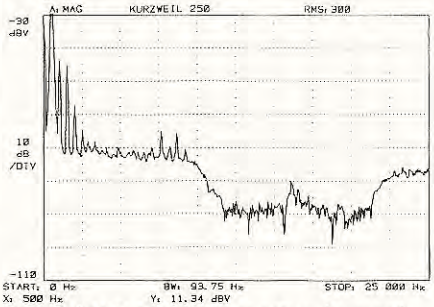

Pitch-Shift Distortion: 997Hz an octave downward.

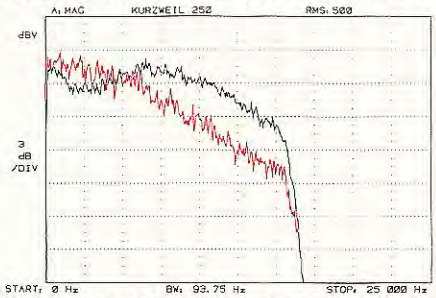

Frequency response: Input-plus-output (black), output-stage (red).

OBERHEIM-ECC DPX-1

While it isn't itself capable of sampling, the DPX-1 plays samples recorded by the E-mu EII, Ensoniq Mirage, Akai S900, and Sequential P2000/2002 (plus Optical Media's CD-ROM library). It can handle samples recorded at rates up to 80kHz (ported in via MIDI); Oberheim's Jim Letts singles out 64kHz as the best-case rate. Few musicians, however, are likely to run into these situations. Since the P2002's top sampling rate, 41,677Hz, is the highest of any of the instruments listed above, we tested the DPX-1 using Softsynth files generated at that rate. Consequently, our DPX-1 test results are most comparable to those of the P2002. It's possible that, in this light, the DPX-1 looks less impressive than it would if it were stacked up against other instruments whose sample formats it reads.

Frequency Response. The DPX-1's frequency response consists of a continuous 45-degree rolloff above around 2.5kHz. The 6dB-down point is at around 8.5kHz—one of the narrower bandwidths we encountered.

Dynamic Range. Aside from its high output level, the DPX-1 ranks about average in all of our dynamic range tests. The "input-plus-output" noise floor plot, in which the DPX-1 played the P2002's silent sample, shows low-frequency energy that doesn't appear in the P2002's plot of the same test.

Distortion. In the 997Hz test, the DPX-1 measured more distortion than any instrument except the Ensoniq Mirage. This is due largely to a third harmonic that reaches to 33dB below the fundamental. (Other low-order harmonic components appear in the 997Hz plot as well.) The 10,007Hz plot, whose corresponding THD+n is less than that of the P2002 in the same test, reveals some inharmonic components and a general noise floor that rises by as much as 5dB at the bottom of the frequency range. On average, the DPX-1's output-stage THD+n (weighted) equals that of the P2002.

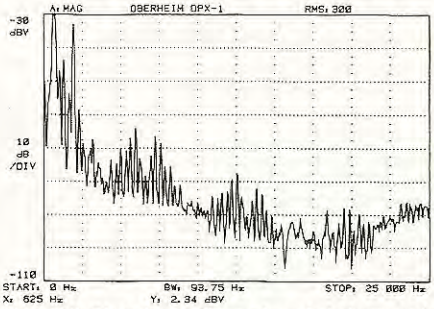

Pitch-Shift Distortion: 997Hz a minor sixth downward.

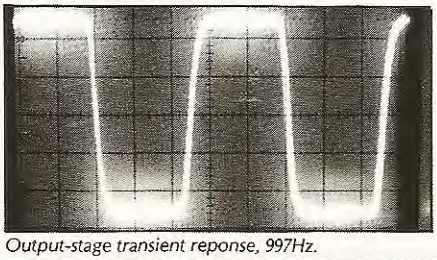
Output-stage transient reponse, 997Hz.

Pitch-Shift Distortion. At different transpositions, the DPX-1's pitch-shift scheme adds different kinds and amounts of distortion, mostly inharmonic. One helpful effect of its limited frequency response is that imaging components transposed down into the audio range aren't nearly as prominent as similar components produced by the P2002.

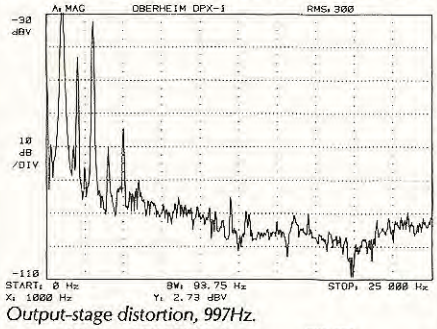

Output-stage distortion, 997Hz.

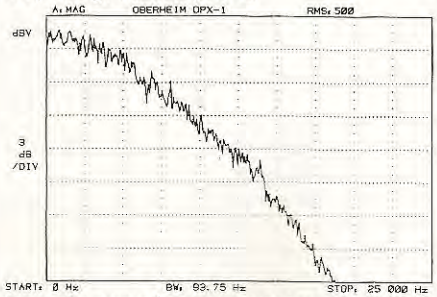

Input-plus-output frequency response.

(Again, this is how imaging appears *in the digital domain.* A proper reconstruction filter is capable of filtering out the distortion, transforming the sample file represented above into a perfectly rounded analog sine wave.)

In order to avoid interfering with these modulation patterns, we used samples of one second or longer, and made every effort to keep our loops from interfering with the secondary cycles of amplitude modulation.

THANK HEAVEN FOR STANDARDS

THE SAMPLERS THEMSELVES DIDN'T make things any easier. For instance, the Kurzweil 250, which dumps samples not via MIDI but through its own QLS format, won't transfer sample data at all if it sees less than 50 milliseconds worth of individual samples before the loop starts. What's more, when sample files in QLS format are translated by a Kurzweil software utility into a format readable by Digidesign's Sound Designer, which we used for making loops, Sound Designer's loop-point markers *disappear* from the computer screen as you try to edit them. Astonishing.

When importing sample files from Alchemy, both the E-mu EIII's and the Akai S1000's sample dump standard implementations inexplicably move start and end points by only an individual sample or two—just enough to introduce an annoying click. If the goal is to produce test results rather than music, one might expect to be able to live with a clicking loop. Unfortunately, the H-P 3561A isn't quite so forgiving. The slightest hint of a click results in a distinctive spectrum display in which the fundamental's peak spreads out across a high noise floor, rather than standing straight and tall. The composite plot below shows a clicking loop progressively causes the fundamental to spread:

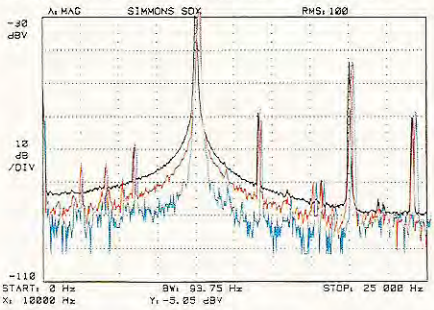

The analyzer reacted this way whenever a looping problem was present, or whenever an instrument's output itself contained some kind of clicking. When the clicking couldn't be eliminated, it was usually possible to get an acceptable spectrum plot by limiting the number of averages taken by the analyzer. This is why some of our distortion plots show RMS numbers of less than 300.

THE 10,000Hz QUESTION

ACCORDING TO THE ORIGINAL PLAN (drawn up in consultation with Digidesign's hardware specialist, Dave Clementson), our distortion tests at all three frequencies were to be conducted using full-code sine waves. (Softsynth automatically computes sine waves at 100% amplitude, and when sampling test signals for input-plus-output-stage tests, we always managed to get within 90% of full code.) The notion of testing 10,007Hz at 100% amplitude didn't sit well with many of the engineers with whom we consulted. They pointed out that the amount of energy at around 10kHz in most music is very small. In the words of Mark Tsuruta and Marty Frazu of Roland's technical staff, "Few, if any, musical instruments have fundamentals at 10kHz. Most musical instruments have fundamentals around 200Hz to 2kHz. The important harmonics that define instrument character are in the 2kHz to 7kHz range."

As a factual statement, this is undeniable. However, the importance of audio frequencies above 7kHz shouldn't be underestimated, particularly by players who use their instruments for sounds other than your garden-variety strings and brass. We reached this conclusion after running a couple of admittedly unscientific, but nonetheless telling, experiments. Connecting the output of the office CD player to the spectrum analyzer, we played a Biograph release of a Scott Joplin solo piano rag. The resulting plot shows that, indeed, there is a rather pitiful amount of energy at 10kHz in solo piano music.

But when we threw on a CD of Bon Jovi ripping through "Shot Through The Heart," the spectrum analyzer told a different story. In the Bon Jovi plot, 10kHz is only 10dB down

ROLAND S-550

The Roland S-550 is one of the big surprises of the survey. Its average distortion rating is lower than that of every other sampler we tested, including the Akai S1000. Unfortunately, it doesn't bear out its promise when it comes to dynamic range and, especially, bandwidth. Our guess is that Roland could work a little engineering magic and come out with *the* killer mid-priced sampler.

Frequency Response. The S-550's 30kHz sampling rate demands steep filtering somewhere below 15kHz, and there it is. The input-plus-output frequency response falls gently from 5kHz or so to the cutoff point, at 13.5kHz. The output-stage curve exhibits an emphasis between 0Hz and 3kHz, and a 2dB notch between 3kHz and 6kHz. The notch, smack-dab in the middle of the area of highest aural sensitivity, will affect any sample ported into the S-550 for playback.

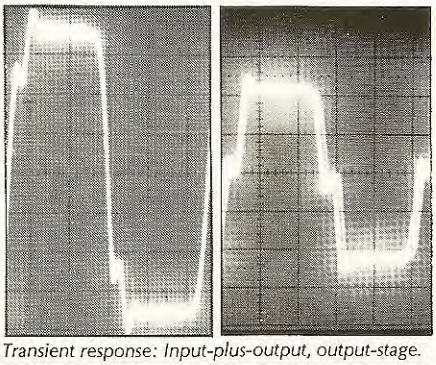

Transient response: Input-plus-output, output-stage.

Dynamic Range. The S-550 has one of the hottest output levels we saw, and its noise floor measurements should be regarded in that light. The quiescent noise floor is below average, but the S-550 sits squarely in the middle of the pack for both output-stage and input-plus-output residual noise. Incidentally, the noise floor plots show a prominent spike at 15,570Hz, the frequency of video signals. It probably represents leakage from the S-550's video display port.

Distortion. As we've mentioned, distortion specs are the S-550's forte. The weighted THD+n measurements, less than one tenth of a percent in all cases, are slightly higher for the input and output stages together than for the output stage alone. The output-stage plots don't look as squeaky-clean as we would expect, revealing a multitude of low-level inharmonic spikes. The output filter's steep cutoff appears as a trench in the noise floor of nearly all distortion plots.

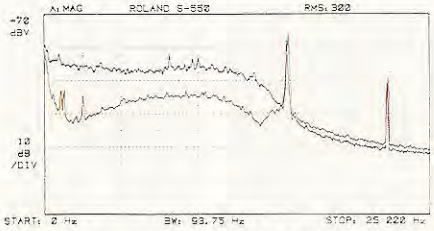

Residual noise: Input-plus-output (black), output-stage and quiescent (red).

Pitch-Shift Distortion. Using a fixed-rate interpolation scheme, the S-550 transposes beautifully, maintaining remarkably low and consistent distortion levels throughout the intervals tested. (What distortion there is is inharmonic.) The imaging components produced by oversampling and fixed-clock-rate techniques are nowhere to be found.

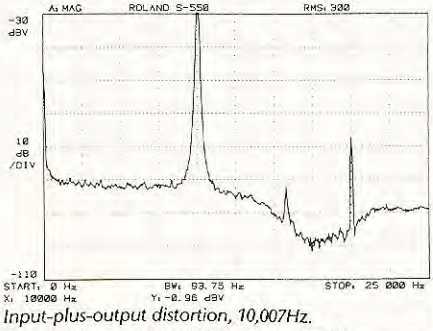

Input-plus-output distortion, 10,007Hz.

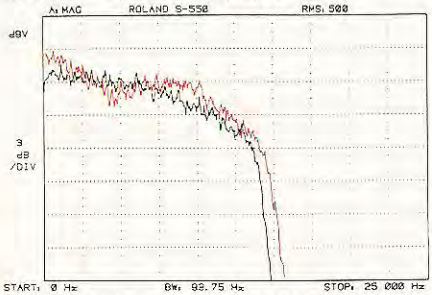

Frequency response: Input-plus-output (black), output-stage (red).

from 5kHz, one of the most important areas of the audio range.

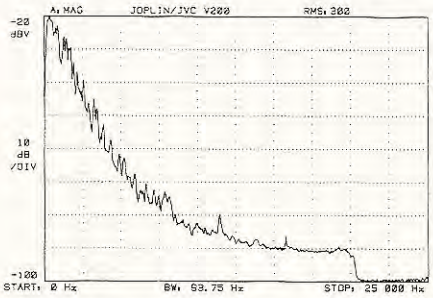

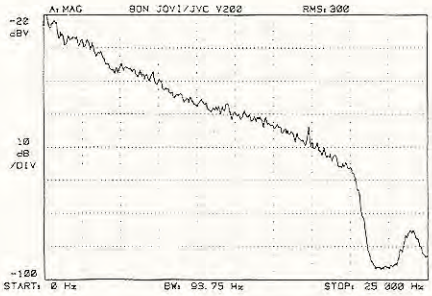

(Incidentally, note the remarkable similarity between the Bon Jovi plot and plots of pure noise which constitute our frequency response test. We hesitate to comment on the implications, but the resemblence is undeniable.) Likewise, analyzing an isolated splash cymbal from the Sound Sampling Collection, Vol. 1 CD revealed that 10kHz is only about 10dB down from the loudest frequency.

At first, our distortion test results seemed to indicate that Roland's point of view is reflected in the design of the S-550; a full-code 10,007Hz tone is about 20% distortion (weighted) by the time it leaves the instrument's output jack. But when similarly exaggerated distortion showed up in such high-end instruments as the Akai S1000 (around 6%, weighted) and the E-mu EIII (around 5%, weighted), we began to wonder if we were making a mistake.

Dave Rossum, chief of research and development at E-mu, suggested that we experiment by lowering the amplitude of our 10,007Hz test signal by a single decibel. We tried it and, lo and behold, most of the distortion disappeared from the outputs of the problem instruments. The decibel scale being logarithmic, 1dB down represents a 10% decrease in signal amplitude. As hot as it may appear, the level of 10kHz in the Bon Jovi CD, -10dB, is equivalent to a mere 33% of full code. Desiring to give the samplers as rough a workout as they might possibly encounter under real-world conditions and at the same time be fair, we adjusted our 10,007Hz test tones to -3dB, or 72% of full code. We kept 97Hz and 997Hz at 100%.

DYNAMIC RANGE

THE DIFFERENCE BETWEEN THE LOUDest and the quietest signal an instrument can record or reproduce is called its dynamic range. According to digital audio theory, there's a formula that describes the dynamic range available from a particular bit resolution: Dynamic range equals 6.021 times the number of bits. According to this formula, a 16-bit system yields 96dB of dynamic range, while 12 bits yields only 72dB. This is a theoretical best case. Whether or not it holds true in the real world depends on the effectiveness of the electronics that surround the analog/digital conversion process.

About three quarters of the way through the test process, we learned how to measure dynamic range in terms of the formula. The recipe goes something like this: Program, from scratch, a set of 997Hz digital "sine waves," in Sound Designer format, whose amplitudes span a single increment at each of the bit resolutions and sampling rates required. Port file to sampler. Play back sampler. Subtract output level when playing back this file—which generates the lowest level the instrument can produce while playing a sample—from the output level when playing back a full-code 997Hz sine wave sample. The resulting figure is the sampler's true signal-to-noise ratio. Repeat until all samplers are tested.

This would have required more custom computer programming than we were prepared to deal with, and besides, we had already devised a set of dynamic range tests. Considering 0dB to be an instrument's output level when playing a full-code sine wave at 997Hz (with the volume control full up), we measured each instrument's output level under three conditions. (1) Quiescent noise

Continued on page 138

SEQUENTIAL PROPHET 2002

One of the distinctive aspects of the Sequential Prophet 2002 is consistency between output-stage and input-plus-output performance. The Turbosynth square wave comes out remarkably square, and the sampled square wave, while rounded, is unusually well preserved. Sampled silence plays back as quietly as pure digital silence. The two sets of distortion plots are virtually identical. The only difference is in frequency response, where the reconstruction filter is lax. Despite high distortion levels, these characteristics suggest that the P2002 is a good choice for those who record their own samples.

Frequency Response. A steady rolloff makes the P2002's input-plus-output bandwidth (13kHz or so) more limited than we would expect from a sampling rate of 47,677Hz. The output-stage response, on the other hand, challenges Dr. Nyquist. Less than 40dB down at 20kHz, it prompts us to suspect that the difference between the unweighted and weighted THD+n figures at 10,007Hz (5% versus 1%) is attributable to supersonic images.

Dynamic Range. The P2002 is quiet for a 12-bit sampler. Quiescent noise and output-stage residual noise are average. The input-plus-output noise level, though, is just behind the 16-bit Akai S1000, E-mu EIII, and Simmons SDX (and ahead of the 16-bit Casio FZ-10M).

Distortion. The P2002's distortion performance is decidedly inferior to most of the samplers tested. The 997Hz plot is messier than either of the others, although at 10,007Hz, the relative noise floor is high due to the rolloff in frequency response. The distortion at 997Hz is basically harmonic, so it probably doesn't sound as bad as it looks. The distortion components at 10,007Hz look like more of a problem.

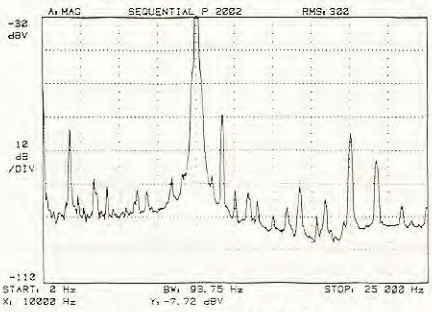

Input-plus-output distortion, 10,007Hz.

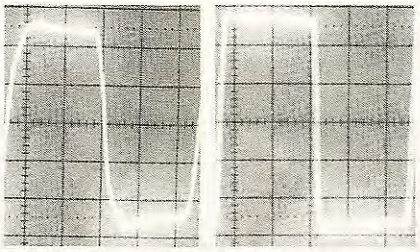

Transient response: Input-plus-output, output-stage.

Pitch-Shift Distortion. The P2002's variable-rate clock keeps this well under control. (At the 41.67kHz sampling rate, upward transposition is limited to a semitone). Extreme downward transposition, however, shifts some nasty-looking images into the audible range—a consequence of the reconstruction filter's gentle slope.

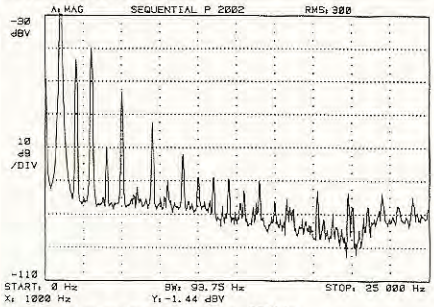

Input-plus-output distortion, 997Hz.

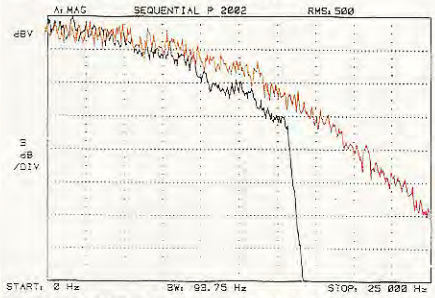

Frequency response: Input-plus-output (black), output-stage (red).

SIMMONS SDX

Designed with drummers in mind, the SDX's layering, editing, and sequencing facilities make it an intriguing sampler for keyboard players. Unfortunately, it falls short in a few crucial areas. For starters, distortion is surprisingly high. (We encountered slight discrepancies between distortion in edit mode and performance mode, but Jim Lindop, head of R&D at Simmons, wasn't able to duplicate this phenomenon.) In addition, we had difficulty creating click-free loops, and periodic clicking appears when the instrument plays back high frequencies, as in our 10,007Hz and pitch-shift tests. Given the SDX's otherwise solid concept and implementation, and Simmons' reputation for quality, we expect that most of these shortcomings will disappear a few software revs down the line.

Frequency Response. Extending to a sharp cutoff at 20kHz, the SDX's input-plus-output frequency response is very broad and reasonably flat. The output-stage response, only 10dB down at 25kHz, leaves the instrument wide open to imaging above the Nyquist limit. The swatch of noise seen above the cutoff point in the input-plus-output plot, similar to that produced by the E-mu EIII in the same test, is probably imaging missed by the reconstruction filter. That it can be transposed into the audio range is verified by the pitch-shift plots.

Dynamic Range. Despite lackluster output and quiescent noise levels, the SDX's dynamic range is impressive. When playing back both zero-amplitude words and sampled silence, the SDX ranks third. The noise floor plots, otherwise smooth and flat, are marred by a spike at 15,750Hz—the frequency of the built-in video monitor's signal.

Distortion. Given its high-end bearing, the SDX's distortion specs are disappointing. Averaged, the weighted THD+n figures are comparable to those of the FZ-10M; individually, they stack up a little worse. Like the FZ-10M, the SDX's distortion is lower at 97Hz and 997Hz than at 10,007Hz. However, it rates significantly better than the FZ-10M's at the higher frequency, and worse in the lower range, where it probably matters more. The 997Hz plots are carpeted with harmonic distortion, and the 10,007Hz plots dominated by a few harmonic and inharmonic peaks.

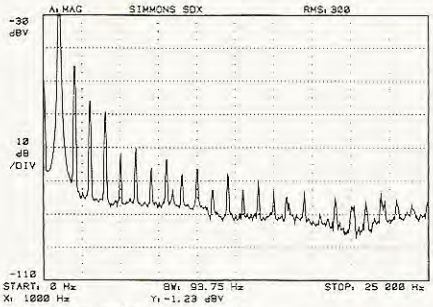

Output-stage distortion, 997Hz.

Pitch-Shift Distortion. We've already mentioned the clicking that accompanies high-frequency playback, which made it difficult to obtain representative plots of the upward transpositions. Otherwise, the plots show that the SDX's variable-rate scheme maintains the original sample's distortion level and composition, except when imaging is shifted down into the audio range.

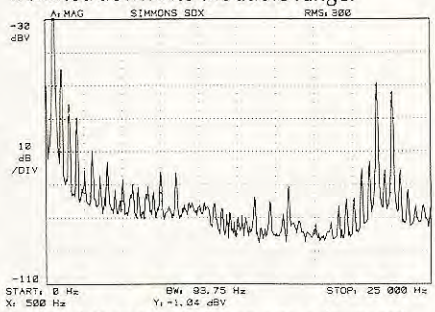

Pitch-Shift Distortion: 997Hz an octave downward.

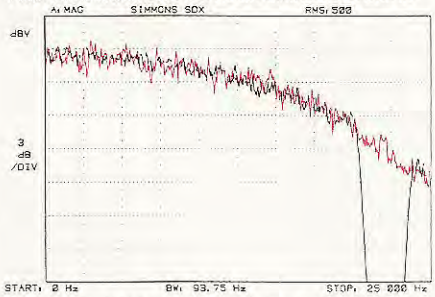

Frequency response: Input-plus-output (black), output-stage (red).

YAMAHA TX16W

Like the Roland S-550, the Yamaha TX16W emerges from our tests with a bunch of monster specs. Also like the S-550, some aspects of its performance are so compromised that the big picture is a good deal less impressive.

Frequency Response. Consider the input-plus-output frequency response curve. Although the anti-aliasing filter's cutoff point is 19kHz, the overall slope produces an effective bandwidth of 11.5kHz. This renders the benefits of a 50kHz sampling rate meaningless.

Dynamic Range. The TX16W makes up for its low output level with good-to-middling residual noise levels. Quiescent noise is about average. Output-stage residual noise is better than any other instrument tested. With the addition of the input stage, the noise level rises enough to bump the TX16W down to the middle of the pack, which is odd, considering that it's an all-digital instrument.

Distortion. Here's where the TX16W shines, outperforming everything except the Akai S1000 and the Roland S-550. The input-plus-output plots are distinguishable from those of the original test signals only by slightly higher noise floors (a result of the frequency response curve) and some low-level harmonic distortion at 997Hz. The 10,007Hz input-plus-output plot is quite amazing, showing no additional distortion components whatsoever. The output-stage plots are equally impressive.

Pitch-Shift Distortion. A textbook example of what's wrong with linear interpolation: Octave shifts look good, but the other intervals tested skew the sample in various ways and by various degrees. Note that, in the graph on page 54, the TX16W is the only instrument whose average change in the amount of distortion induced by transposition exceeds the original amount of distortion in the sample. The benefit of this approach is that imaging isn't a factor in downward transpositions.

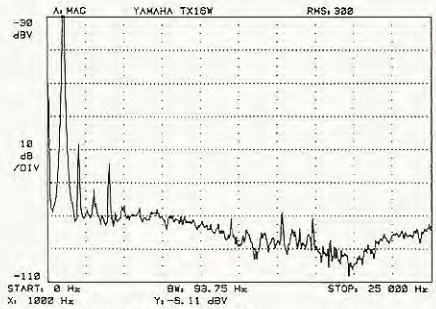

Output-stage distortion, 997Hz.

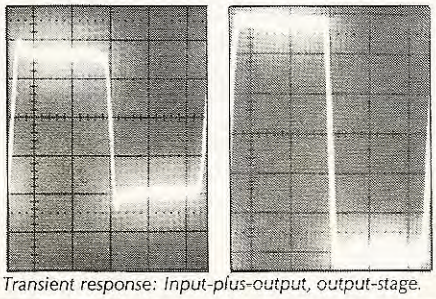
Transient response: Input-plus-output, output-stage.

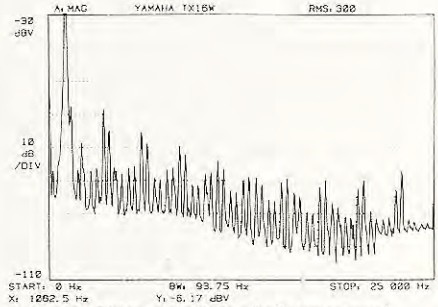

Pitch-Shift Distortion: 997Hz a half-step upward.

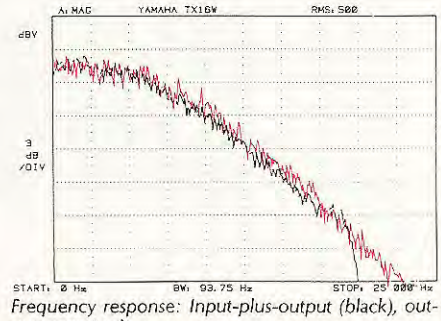

Frequency response: Input-plus-output (black), output-stage (red).

SAMPLER TESTS

Continued from page 64

floor: This is the output level when the instrument is completely idle, with its volume control full up. (2) Output-stage residual noise: This is the level when the instrument plays back a sample consisting entirely of zero-amplitude words. (3) Residual noise of the combination of input and output stages: This is the level when an instrument plays back sampled silence. (Silence was sampled with the input volume about halfway up and nothing plugged into the sampler's input.) These measurements are expressed as a negative number, or so many dB down from the full-code output level. We also noted the actual 997Hz full-code output level, expressed in dBm.

The quiescent noise floor shows, in relative terms, how much noise the mere presence of a given instrument will add to your audio system. A high quiescent noise floor paired with a low 997Hz full-code output level is a bad sign; this means that you'll have to boost an already-high noise level even higher at your mixer to bring the output up to a good listening level. At the same time, you'll also be boosting the mixer's inherent noise, making for a potentially nasty total noise floor. A glance at the dynamic range graph on page 53 confirms that the Casio FZ-10M is the worst offender in this regard.

Playing back computer-generated zero-amplitude words—pure digital silence—exercises only the instrument's output stage. This includes any analog processing modules, the reconstruction filter, and the output preamp. A high-quality sample ported in via MIDI and played back can't be any quieter than this level, even at the moment it decays to silence.

Playing back sampled silence allows us to measure the noise generated by total instrument in operation (input plus output stages), including the anti-aliasing filter, A/D converter, input and output preamps, analog processing modules, DAC, reconstruction filter, and so forth. If you record your own samples, this is the level your instrument plays at the moment a sample decays to silence.

In addition to taking measurements in dB, we plotted the spectra of the three noise floors. Those of the Oberhiem DPX-1 Sample Player are shown, superimposed, below. The bottom line is the quiescent noise floor, the middle line is the output-stage noise floor, and the top line is the noise floor generated by the playback of sampled silence. (The DPX-1 is a playback-only machine, so silence was sampled using the Sequential P2002. Thus, the top line represents the residual noise of the DPX-1's output stage and the P2002's input stage.)

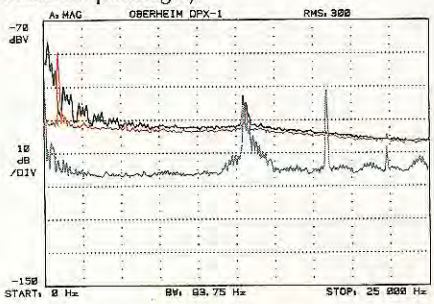

When evaluating noise floor plots, bear in mind that corresponding measurements found in the chart on page 53 represent averages across the frequency range, while the plots measure each frequency bin individually. That's why the plots look much lower than the measurements seem to indicate. Also, the measurements represent a number of dB below the full-code 997Hz output level, while the plots show an absolute scale in dBV. You can't tell how quiet a given instrument's plotted noise floor really is until you've referenced it to the same instrument's 997Hz level. The dB measurements for various instruments, on the other hand, are directly comparable.

PITCH-SHIFT DISTORTION

ONE MAJOR DIFFERENCE BETWEEN samplers and other kinds of digital recorders is that samplers play back recordings at pitches other than the original. We transposed our 997Hz output-stage test tone by a half-step, a minor sixth, an octave, and a sharp 11th in either direction, taking both spectrum plots and THD+n measurements along the way. These intervals produce a variety of ratios with respect to the original pitch, making them a good bet for putting a pitch-shift scheme through its paces. (Some instruments transpose only over a limited range when playing back samples recorded at the highest available sampling rate. Such limitations are denoted in the chart by "n/a.")

It makes sense to us to evaluate pitch-shift distortion on the basis of whether or not it's consistent at all intervals. (Of course, intervals at which we didn't test may be more or less distorted than those whose distortion percentages appear in the THD+n chart.) Since weighting can disguise consistency among the THD+n measurements, the unweighted figures are best for this purpose. Generally speaking, consistent pitch-shift distortion is associated with samplers that use a variable clock rate during playback, such as the Simmons SDX, E-mu EIII, and Kurzweil 250. Fixed-rate instruments, including the Yamaha TX16W, E-mu Emax, and Ensoniq Mirage, are more likely to introduce different amounts of distortion at different transpositions. (For an explanation of the differences between fixed-rate and variable-rate technologies, see page 40.)

While it's relatively easy to spot trends in the pitch-shift distortion measurements, interpreting them presents a couple of problems. For one thing, they all begin with the amount of distortion already introduced by each machine's output stage at 997Hz. Therefore, a simple average of the THD+n at all transpositions would produce a ranking more or less identical to one based on the original 997Hz THD+n measurements. The other problem involves the weighted figures. When transposed, a sample's THD+n will decrease as more of the distortion components are shifted beyond the range of the band-limiting and A-weighting curves. If you consider the adjusted THD+n measurements, some of the most distortion-prone machines show, on average, the least pitch-shift distortion—usually "negative" levels of distortion, in fact, relative to the original 997Hz percentage.

We've tried to make sense of this in the scatter graph on page 54. It plots the original output-stage THD+n at 997Hz as well as the average difference between the original THD+n and the THD+n at various transpositions (all figures weighted). This keeps the relationship between the baseline distortion level and the average difference induced by transposition fairly obvious, so you won't be hoodwinked by the Ensoniq Mirage's low pitch-shift distortion average into thinking that it's a low-distortion sampler. "Negative" distortion doesn't mean less distortion. It's simply a byproduct of weighting. It seems to us that, regardless of whether they fall above or below, the closer the white dots are to the zero-distortion line (the x-axis), the more effective the sampler's pitch-shift scheme.

TESTING TRANSIENT RESPONSE

TRANSIENT RESPONSE REFERS TO A sampler's ability to record and reproduce changes in a sound that happen very quickly. For instance, when you strike a tom-tom, the resulting sound goes through all kinds of contortions during the first fraction of a second (the attack), after which it settles down into an organized pattern (the after-ring). In order to record the tom-tom with reasonable fidelity, a sampler must be able to capture all of the changes occurring during the attack. Its input circuitry must be fast enough to pass along new voltages to the analog-to-digital converter accurately, and the converter must be fast enough to see all of the changes as they occur. Ditto for playback; the digital-to-analog converter must be able to fire off enough values to re-create the initial sonic variations, and the output circuitry must be able to handle the changes without smoothing their rough edges. Transient response is related to frequency response, since a high frequency is nothing more than a transient that repeats again and again.

We tested transient response by sampling a square wave, which consists of instantaneous alternations between full positive and full negative—about as transient as you can get. What came back out of the samplers was rarely square, but this doesn't necessarily indicate a problem. Because transient response is allied with frequency response, we would expect a sampler's anti-aliasing and reconstruction filters—necessary components of a digital audio system—to have some effect on square wave's shape. This is the most likely cause of the "ringing" commonly seen at the top and bottom of the cycle. Digital audio experts caution that the relationship between a square wave's shape and its sound is extremely difficult to determine.

Therefore, rather than attempt to interpret them, we've simply reproduced the results of the transient response test when they look interesting. For comparison, here's an oscilloscope display of the signal used in our input-plus-output-stage transient response tests:

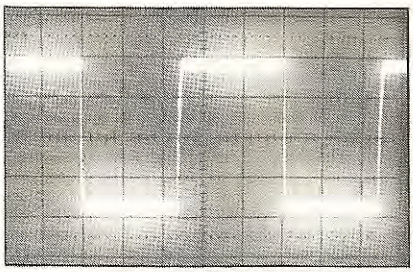

Our "Hip To Be Square" award goes to the Kurzweil 250, which managed to mangle the test signal into the pitiful waveform reproduced below:

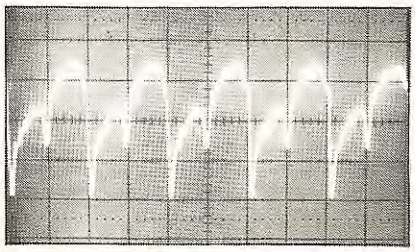

WHAT WE DIDN'T TEST

THIS SURVEY FOCUSES ON SAMPLERS that function primarily as polyphonic keyboard- or MIDI-controlled instruments, leaving out sampling drum machines, sampling delay units, monophonic one-shot samplers, and direct-to-RAM and -disk digital recorders. Many of the instruments we look at are available in both rack-mount and keyboard configurations; in such cases, we tested only the rack-mount model.

Many readers will note the absence of the Fairlight CMI, the New England Digital Synclavier, and the WaveFrame AudioFrame. At the time of our tests, Fairlight was undergoing a corporate reorganization and found themselves too preoccupied to respond to our requests for a CMI. New England Digital was working on a hardware revision which promises to improve the Synclavier's specs, and weren't able to provide an updated instrument soon enough. WaveFrame, whose AudioFrame is still largely unfamiliar to the musical community, felt that their instrument would be better represented in a full review rather than a spec test, and declined to participate.

We actually did test the Marion Systems MS-9C upgrade for the Akai S900. The results had been tabulated and analyzed when a trail of odd circumstances led us to the conclusion that the Marion upgrade had been tested in 12-bit, rather than 16-bit, mode, thus depriving it of the opportunity to show its best stuff. How this could have happened is something of a mystery, since designer Tom Oberheim himself set the unit's internal jumper to boot in 16-bit mode. However, Oberheim confirms that one in a hundred times, when you turn the S900 off and immediately on again, the upgrade will toggle between modes. It appears that we just happened to hit that lucky hundredth time when we rebooted the unit in order to clear out its memory. All is not lost, however. Although the 12-bit Marion results (which slightly bettered those of the stock S900, by the way) were yanked from this issue, 16-bit results will appear as part of a Keyboard Report next month.

BLINDED BY SCIENCE

ONE FINAL *CAVEAT*. THE GENERALLY accepted notion of science seems to be that it consists of cut-and-dried facts. Actually, science consists of interpretations regarding which of the infinite number of cut-and-dried facts in the universe are the important ones. By accepting some facts and discarding others, paleontologists regularly extrapolate the ascendency and extinction of entire species from little more than a few jawbone fragments the size of your thumbnail. We believe that our efforts to produce scientific measurements of sampler performance stand on firmer ground—but it's only a matter of degree.

This is a roundabout way of saying that, while we have a strong conviction that the results presented here are accurate and reproducible, it must be emphasized that the number of variables involved is enormous. Controlling them requires ignoring some facts in favor of others, and some of them are simply beyond control.

For instance, sometimes our noise floor plots changed literally from minute to minute; ripples would would come and go, frequency spikes would appear and disappear. We tried to eliminate this phenomenon by conditioning the office's AC power using a pair of Tripp Lite LC-1800 Line Stabilizer/Conditioners, but we couldn't get it to go away.

Speaking of noise floors, those of many instruments improve drastically after the machine has been allowed to warm up for an hour or so. Distortion levels tend to be more stable, but we did discover that the Akai S950, at least, gives cleaner THD+n measurements when it's cold. (The difference is barely visible in the spectrum plots.) The results presented here represent instrument performance after several hours of warm-up.

Circuit resistance can also affect a sampler's performance. All measurements were taken across a resistance of 50k ohms. Using a 600-ohm load—a conventional condition for testing some kinds of audio gear, but one which few musicians are likely to encounter— made some instruments perform poorly. Was it right to use 50k ohms? All we can say is that many—though not all—of the engineers with whom we consulted agreed that 600 ohms was unreasonable, and that 50k ohms was more realistic.

And then there's the issue of our familiarity with the instruments at hand. We made every effort to open every filter, boost every gain stage, disable every envelope, and so forth, but we may have missed something. The Akai S900 has a "loudness" parameter in its sample edit page that improves signal-to-noise ratio significantly. Discovering it sent us scurrying over to the S1000. There's nothing analogous in the S1000's sample edit page, but eventually we found an extra gain stage, along with a filter cutoff boost, hidden in one of the keygroup pages. This meant testing the S1000 over again.

This is just the tip of the iceberg.

Issues such as these strongly suggest that comparing specs produced in different test environments is a mistake. Our results, whether or not they agree with those you may find published elsewhere, at least share the virtue of having been generated under a consistent set of test conditions. ■