Inside the

To those of us who make a living playing or recording music, the advent of the new computer-controlled synthesisers is both exciting and frightening. Exciting because the new machines open up entirely new areas of sound generation and control, give us access to powerful new compositional and arranging tools, and promise to make our lives easier by eliminating so many of the tedious tasks always associated with making music and putting it down on tape. Frightening because the technology will soon be so prevalent that those who are not comfortable with it, or lack the skills to take advantage of it, may well find themselves out of work.

Among the most interesting and versatile of the new machines is the Synclavier (formally known as the *Synclavier II*—the company's first model was discontinued in 1979). The Synclavier (pronounced, for the record, 'SYN-kla-*veer*') is at once a synthesiser, a recorder, a transcriber, and a tool for sophisticated computer-controlled composition using either real-time inputs or off-line programming.

The *Synclavier* developed out of a research grant to a group of musicians and computer scientists working at Dartmouth College, in Hanover, New Hampshire. The team's first project was to design a digital synthesiser module, several of which could be hooked into Dartmouth College's Data General *Nova* time-sharing system, so that up to 16 students could use the computer simultaneously for individual composition projects. The grant stipulated that if the project was successful, the results should be taken out of the academic sector—ie somebody should make a few bucks out of 'em.

Successful it was, and in 1976 New England Digital was formed and set up shop just across the Connecticut River in Norwich, Vermont. Two years ago, the company moved downstream into bigger headquarters at White River Junction. Today, the company manufactures, assembles and tests the *Synclavier* in a converted plumbing-supply house, employing some 35 full-time staff.

Although it's not always readily visible ("In 1980 it was hard to sell musicians on the concept of a computer," says software development director Jeff Risberg), at the heart of the machine is a proprietary 16-bit 64K rack-mountable minicomputer, designed expressly for real-time control and manipulation. According to company president Sydney Alonso, the computer used in today's machines is essentially the same as the Dartmouth team's first model—although the software has developed quite a bit. (The computer has other uses as well: about 30% of New England Digital's sales are non-musical, mostly in the fields of laboratory data gathering, manufacturing automation and communications networks.)

In the basic *Synclavier*, the computer is controlled by a 5-octave piano-type keyboard, above which is a panel containing 128 lighted buttons, a single control knob, and a 4-digit LED readout. Storage is handled by a double-density 5¼ in disk drive. Optional disk configurations include 8 in floppies and 5¼ in and 8 in Winchester hard disks and the computer itself is expandable, in increments of 32K. All of the components (except the larger disk drives) can be packed into custom-designed ATS hardshell travel cases.

Making waves

Inside the computer are the oscillators—digital-to-analogue converters which are also a proprietary design. Each synthesiser 'voice' starts with a pair of oscillators. The voices are available in hardware packages of eight pairs—as many as 128 voices can be installed, although the usual configuration is either 16 or 32 voices.

The oscillator pairs use a combination of additive and FM synthesis. One oscillator is the 'carrier'. It sets up a sine wave, over which 24 harmonics can be generated and balanced. The individual harmonics are called up—singly or in

Paul D Lehrman

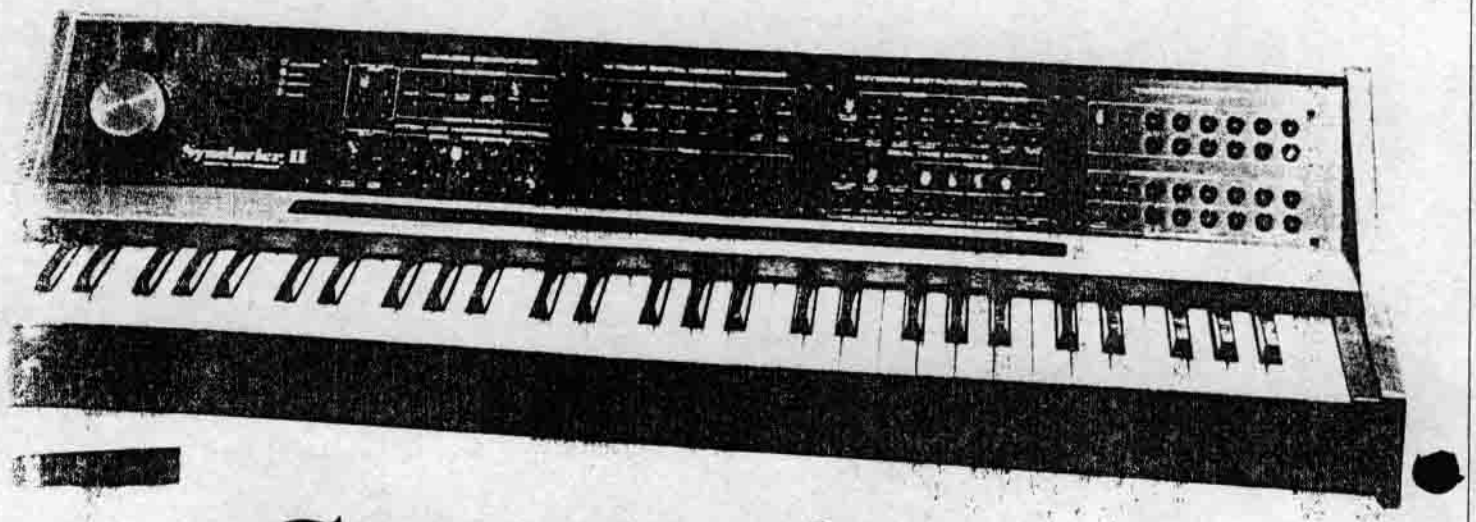

Synclavier

groups—from a section of the panel buttons, and are adjusted with the control knob. As each is adjusted, its relative level shows up on the readout. The completed carrier wave is then given an envelope (called, not surprisingly, the 'volume envelope') with six parameters: delay, attack rate, peak level, initial decay rate, sustain level, and final decay rate; all of which, again, are controlled with the knob and read on the display.

In fact, on the basic *Synclavier*, the control knob and the LED readout are the only way of adjusting and examining the many synthesiser parameters—the last button pushed dictates which parameter is being addressed. (Pitch, of course, is determined by the keyboard.) Four single LEDs next to the readout indicate whether the units for the parameter chosen are Hz, ms, dB, or are on an arbitrary scale.

The second oscillator in the pair is an FM modulator. It, too, starts as a sine wave (and stays that way), whose frequency is determined by a setting called 'ratio'. A 1,000 Hz carrier modulated with a sine whose ratio is 0.1, for example, will produce sidebands to the carrier that are 100 Hz apart. (The modulator can also be set to a constant frequency between 0.1 and 999 Hz, independent of the carrier frequency.) This creates non-harmonic tones, which have the potential to be much more interesting than the harmonic ones created by the additive synthesis process.

The modulator wave is then given its own six-parameter envelope: the 'harmonic envelope'. The higher the modulator wave's amplitude, the more sidebands are created (and the greater their level) which makes for a 'denser' sound.

While the volume envelope controls the overall shape of the sound, the harmonic envelope independently controls its density, so that a sound can change from dense to simple, or *vice versa*, and maybe back again, as it plays. (A recent software update allows the amount of FM to be varied across the keyboard, so that lower notes can be made to sound more brilliant, with more sidebands, while higher notes are less 'grainy'.)

Adding effects

The composite waveform can then be assigned vibrato, which is another form of frequency modulation but uses only low frequencies. One of five simple waveforms can be used for vibrato (which can be applied to both carrier and modulator or to the carrier alone) and rate (up to 50 Hz), depth (up to two octaves) and attack time (up to 10 s) can be specified. The vibrato wave can be inverted, biased (so that it only acts to above the played pitch) or quantised, so that instead of altering the pitch smoothly, it jumps in discrete (usually semitone) intervals. 'Tremolo', or amplitude modulation, can also be added, completely independently of the pitch vibrato.

The synthesised wave can also be given portamento (glissando) which can be set to sweep from one end of the keyboard over an interval of as long as 60 s, with either linear or logarithmic motion. The volume envelope can be told to automatically repeat, at a rate of up to 100 Hz, or to attack separately (arpeggiate) each note of a held chord in sequence.

The scale of the keyboard can be adjusted so that each half-step can sound anything between $1/_{64}$ of a semitone(!) and a minor third, and each note within an octave can be tuned separately, allowing just, Pythagorean, mean-tone, or non-Western intonations. In addition, a custom Morley foot pedal can be plugged in to control certain 'real-time effects': overall volume of the synthesiser output, vibrato depth, portamento rate, attack time, decay time and/or peak and sustain level of either or both the volume and harmonic envelopes.

Designing instruments

All of these parameters—carrier and modulator waves, envelopes, vibrato, portamento, tuning and effects—define what New England Digital calls a 'partial timbre'. (The choice of nomenclature is admittedly unfortunate but it persists, according to Jeff Risberg because it exists in some of the system's patent documents.) Up to four partial timbres, each with its own distinct set of parameters and each tunable over the full audio range in intervals of 0.1 Hz, can be combined to create an instrument. Individual partial timbres can be 'chorused'. By duplicating the settings into another pair of oscillators and detuning them—either slightly, for phasing effects, or grossly, for automatic harmonisation—and then the entire instrument can be 'double-chorused'. Parameter settings for one partial timbre can be 'bounced' to another and then altered slightly, which allows the rapid creation of incredibly rich, fluid sounds.

The keyboard can be split into segments as small as one note, each calling up different combinations of partial timbres and the volume of each partial timbre can be set to taper off at the upper and lower edges of the keyboard segment. Each partial timbre can have its own set of split points. This allows authentic reconstructions of instruments, like piano, harp, or strings, whose timbre does not remain consistent throughout their range.

At present, the *Synclavier* keyboard is a simple switching type but a pressure- and velocity-sensitive version should be available soon. Two prototypes are on the workbench at New England Digital: one is a simple weighted lever, while the other uses a fairly sophisticated acoustic-piano-type escapement.

A new option allows sophisticated control of the sounds within a stereo field. Individual partial timbres can be placed at any of 100 left-right locations. Twenty different stereo modes are available, including static, keyboard- or control knob-based panning and automatic panning with selectable waveform and rate.

The stereo option is a good example of how New England Digital can add new features to the

70 ▶

Synclavier without changing the basic control hardware. The panel buttons marked 'vibrato' have three modes: off, on, and 'blinking'. The first two are self-explanatory but in the blinking mode, these buttons control the stereo effects without affecting the vibrato settings. Potentially, all 128 buttons on the panel could have multiple functions. One planned update, which will no doubt be controlled somehow from the panel, will be individual audio outputs for each partial timbre.

Another effect that can be used in real-time performance is pitch bend, which can be defined either with the main control knob or with a ribbon controller. Other foot-pedal jacks allow access to sustain, hold and on/off switching of the portamento, repeat, and arpeggiate functions.

The instruments are stored on disk in groups of eight. Eight 'timbre banks' fit on a 5¼ in floppy. Access to an instrument within a bank is instantaneous—switching banks takes a few seconds for the disk to load it in.

Recording

The amount of control over synthesised sounds in the *Synclavier* is impressive but it is only a small part of the story. Another important feature is the machine's built-in 16-track 'memory recorder'. Each track can play only one instrument but multiple keyboard lines can be overdubbed on a single track without erasing. While notes are being recorded, they can be rhythmically 'justified'—the timing of the beginning of a keystroke can be automatically corrected to sound right on the money of a user-determined rhythmic pulse.

The recorder includes the usual functions associated with multitrack tape recorders, like start, stop, record, play, erase, solo and punch-in and punch-out (operable from the control panel or foot switch) and a few not available with analogue tape: rewind and fast forward do not change the pitch of sound, just the speed, although the two controls can be shuttled, with automatic slow-down and speed-up, just like on a tape machine. There is a loop function with independent start and stop points for each track. Pitch or speed of an individual track can be altered each without affecting the other, or any other track. Notes recorded on one track can be bounced to another (containing a different instrument) erasing the original track or not, allowing automatic orchestral doubling; their start time can be altered in the process. Timbres of already-recorded tracks can be changed without altering the notes.

A digital metronome (click generator), which is independent of the speed control and can read out in Hz, frames, or milliseconds, is included. Information from the pedal is recorded separately (either in real time or after the track has been recorded) so that track-by-track dynamic mixes can be accomplished. Depending on the size of the computer memory, up to about 8,000 notes can be stored. Up to six compositions (of 8,000 notes each) can be recalled instantly from the *Synclavier* control panel, and individual tracks recorded in one sequence can be loaded into another.

Studio Interfacing

Because the number of voices available in the system is limited (128 voices *is* a lot but putting all that hardware into the machine raises the price to well over $100,000) and because 16 oscillator pairs can be eaten up by a single note (if it uses four partial timbres with chorusing and double-chorusing) it makes sense to be able to dump the *Synclavier*'s output to multitrack. This is provided by an external synchronisation system that generates a 50 Hz pseudo square wave, which can be recorded on a tape track and which the *Synclavier*'s recorder can recognise and lock up to on subsequent passes. The sync system can also lock two *Synclaviers* together and can read external pulses generated by timing generators, drum machines, or other synthesisers.

The combination of the digital metronome and the sync function makes the *Synclavier* particularly useful for film and video composers because the metronome can be controlled by the software, it can be set to count frames at any rate, for any format, real or imagined. Speed and timing of individual tracks and complete compositions can be controlled with a very high degree of precision, so music can easily be made to fit the most unwieldly of film edits. There is no full SMPTE interface available as yet but the company is working on it.

In addition, there are several other ways to link up the *Synclavier* with external equipment. Keyboard trigger, gate, and voltage signals appear at separate output jacks, where they can be patched to control other synthesisers or processors. In addition, programmable control voltages that track key motion and are settable from the front panel can be output to external high pass, low pass, or bandwidth filters.

Getting into the computer

Up till this point, we've been talking primarily about your basic *Synclavier II*, costing between about $14,000 and $30,000, depending on the number of onboard oscillators. But there is a wide range of options to the system and new ones are being developed all the time. These options use the identical computer as the basic *Synclavier* (although some of them require extra memory and/or storage) but provide more sophisticated ways of controlling it.

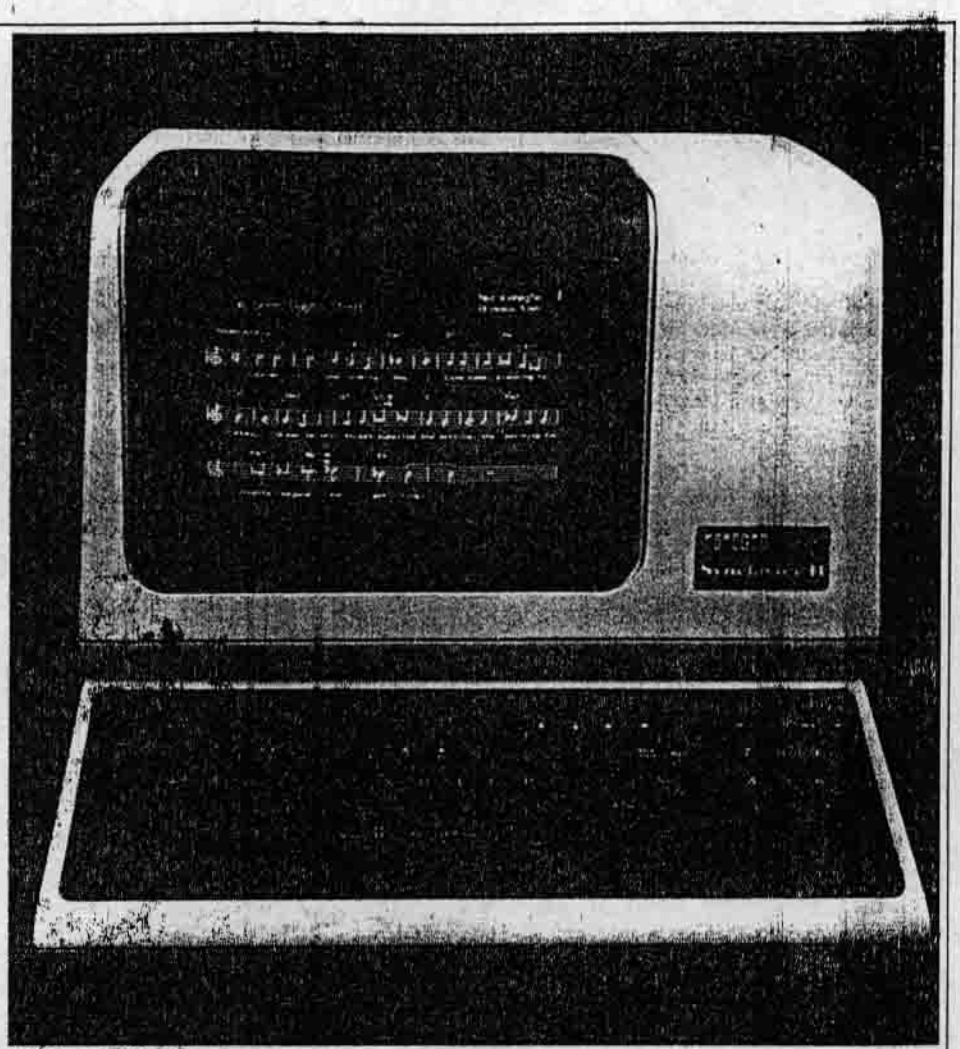

Terminal Support Option

For about ten grand you can get the 'Terminal Support Option' (TSO), the first major add-on that New England Digital developed. About half of the 500-odd *Synclavier II*s sold to date, have left the shop equipped with this option.

The TSO consists of an ASCII keyboard attached to a video display terminal and the appropriate software. It increases the flexibility and speed of the system dramatically. (It also increases the size of the exceptionally well-written manual—from a manageable 140 pages to, when all the software documentation is included, well over 600 pages. Luckily, it all comes in a loose-leaf binder.)

One feature of the TSO is the Timbre Display System. This augments the *Synclavier*'s readout functions by simultaneously displaying multiple parameters. A menu of alphanumeric and graphic formats can show various combinations of envelope, spectral and effects parameters for one or all of the partial timbres in an instrument. Displays are automatically updated when any parameter is changed and new and old settings can be displayed simultaneously. Data can also be output to a printer.

Another very useful section of the software is a composing and editing language called 'Script'. With this program, both note files and timbre control settings can be entered from the computer keyboard. A sequence can be entered, note by note and track by track, in a form that specifies pitch, timing, volume, timbre and articulation for each note, and it can then be edited much like text in a word processor.

72 ▶

Inside the Synclavier

Groups of notes (called 'motifs'—I suppose if the system used a lightpen, you could call them 'lightmotifs'...sorry) can be transposed, inverted and reversed. Loops can be defined, just as on the keyboard control panel and various global functions can be performed on a finished composition, such as tempo, key, volume and articulation.

Perhaps most important: sequences entered into the *Synclavier*'s memory recorder in real time from the piano keyboard can be converted into Script files, so they can be edited off-line—mistakes can be corrected, individual tracks can be 'slipped', and many other useful functions can be performed.

Yet another section of the TSO is 'MAX', a musical programming language based on a subset of XPL, that gives the user direct control of the *Synclavier*'s computer. It is a powerful tool for the serious computer composer who is not necessarily interested in limiting his input control to the piano keyboard and foot pedals—it can perform complex mathematical functions on notes, loops and other functions, and can design infinitely complex envelopes. The range of MAX's functions is a little beyond the scope of this article but suffice it to say it allows the ambitious user to customise quite completely the *Synclavier*'s operation to suit his needs and also allows New England Digital to continuously develop alternative operating systems for the *Synclavier*.

Guitar interface

Another way for musicians other than keyboard players to take advantage of the wonders of digital synthesis is New England Digital's recently-released guitar interface for the *Synclavier*. The interface (which was extremely well demonstrated by Pat Metheney at the last New York AES) works with a Roland *GR* guitar, and costs about $4,500, plus the price of the guitar. It uses the guitar's built-in 'hexaphonic' pickup and preamps to generate six channels of pitch and amplitude information, which are then processed by the interface circuitry at a rate of better than 200 samples/second. Unlike many other guitar synths, the signals from the *Synclavier* guitar interface do not emerge as voltages—rather they are converted into digital codes, which are fed directly to the computer. This method is designed to eliminate drifting and improve pitch-tracking accuracy on transients.

The option comes with a 16-button control panel which attaches to the guitar body. Not surprisingly, the buttons serve multiple purposes: they choose instruments and timbre banks, they select tracks for the memory recorder and control its 'transport' functions and they switch among several control schemes for reading the guitar data having to do with whether, and how, the pitches being read will be quantised. Notes entered on the guitar are treated by the system exactly the same as notes entered on the keyboard and can be used in all of the other *Synclavier* functions, including Script files. A control on the guitar allows real time mixing of straight guitar and synthesised sounds.

Transcribing

Just as drummers have become some of the best customers for drum machines, music copyists are beginning to look into *Synclaviers* as a way of protecting their careers—in particular, due to the Music Printing Option (about $3,000 including a Prism dot-matrix printer, plus the cost of the Terminal Support Option). The printing program uses Script files for input, and the click generator for timing. Under direct control of the composer are such features as time signatures, key signatures, transposition (for non-concert-pitch instruments), choice of any of four clefs, note resolution, whether accidentals are to be treated as sharps or flats, and which tracks in a note file are to printed out. All of these parameters can be changed in mid-piece. (For now, all tracks must share the same note resolution and time signature but a planned software update will allow those parameters to be assigned independently.)

A lot of sophisticated printing routines are available. For instance, ties between notes are optional: in the 'classical' format, off-beat sustained notes are tied, while in the 'jazz' format they are treated as single notes. Triplets, quintuplets and any other odd tempo divisions can be printed accurately. Markings like dynamics, tempo, articulation, titles, instrument names, vocal texts, and page and measure numbers can be added. Vertical and horizontal spacing can be adjusted. Separately-recorded

The author ponders...

tracks can be combined on one staff (with stem directions indicating two different voices, if desired) or printed out in orchestral-score form. The only limitation to the size of the score is the size of the printer paper—but it is not difficult to print scores in sections and then paste them together. Except for a few microscopic jaggies, the output of the dot-matrix printer is just about as good as that produced by a linotype or plotter—and it's much faster.

Playing with real sounds

The *Synclavier* feature that holds the most promise (and for some the biggest threat) is the Sample-to-Disk option—about $10,000 or so on top of the TSO. The system samples with 16-bit words in a PCM format at a rate of up to 50 kHz. Depending on the amount of disk storage available and the sampling rate used, it can theoretically store up to 54 minutes of sound.

A sound to be sampled is entered from a microphone or other analogue source and converted through a D/A into a computer file. The file is displayed on the screen—first in its entirety, so its overall envelope can be viewed and then (in larger form) in smaller time increments, with resolution as high as one sample, to give a close-up view of the individual waveforms recorded. Different segments of the file, of whatever length is desired, can be called up for display, and the vertical (amplitude) scale can be adjusted.

Once the file is in memory, it can be called up on the *Synclavier* keyboard, which will play it at different pitches by changing the playback sampling rate. (If a playback rate higher than 50 kHz is called for, nothing will happen—so you can't play a sound recorded at the highest sampling rate at a higher pitch. To overcome this, the system allows files to be automatically resampled at different rates.) New attack and decay parameters can be impressed on the sound, vibrato can be added and individual segments of files can be extracted, looped or spliced together with exquisite precision so a sound can begin say, as a trumpet, and end up as a human voice.

Again depending on sampling rates and storage space, up to eight files can co-exist in the memory. Many of the keyboard effects can be used on sampled sounds, including split and pitch blend and some of them can be under foot-pedal control.

In addition, there are a wide variety of analysis and filtering functions included in the software. Fast Fourier Transforms can be performed, with resolution of up to 8,192 points. Analysis can take the form of spectral density, or of relative amplitudes of the waveform's component harmonics. All these displays are two-dimensional graphics but a pseudo-3D display mode is available that shows the change of spectral parameters over time. Images on the screen can be saved as hard copy, using a printer interface that matches the screen graphics pixel for pixel.

Filtering in the system is performed at the digital level, which gives a high degree of precision and good phase-linearity but has the drawback of being rather slow—computations of a filtered waveform taking several minutes are not uncommon. The range of filters available is almost unlimited—bandpass, bandstop, high-pass and lowpass, combs, impulse-trains and various kinds of time- and phase-dependent functions. (According to Jeff Risberg, samples recorded at the factory are subjected to a 14 kHz deep-notch filter, to overcome the whine of the air-conditioning system.) The filters can also be designed to model the impulse response of an acoustic space, so that they can add (or remove) room effects to a sound.

For now, the Sample-to-Disk option is homophonic—although several sounds can be accessed from the keyboard, only one at a time can actually sound. Using the various *Synclavier* functions, it is possible to lay down sampled sounds polyphonically, with the assistance of a multitrack tape recorder: synthetic patches that resemble the real sounds (especially in terms of their envelope) can be set up and played into the *Synclavier*'s memory recorder and then each track of the recorder played back individually, using a sampled file for its timbre and dumped (using the sync function) on tape.

As it stands, the Sample-to-Disk option is a fairly sophisticated version of an *Emulator*, although it does have far greater modification capability and control. (Reportedly, a famous American singer who is not known for accurate tracks recorded vocals into a *Synclavier*, so that the record producer could correct them for pitch and timing before dumping them onto the master tape.) But New England Digital hopes this spring to bring out the first of a series of polyphonic sampling options that will do much more.

The first version will support four simultaneous voices, while subsequent updates will be 'N-voice'. Plans are to have the new software to be able to automatically analyse sampled sounds and resynthesise them, which will allow most of the *Synclaviers* control functions—now usable only with internally synthesised sounds—to be applied to real sounds. In addition, sampled sounds and synthesised sounds will be able to be played simultaneously.

At that point (and with the addition of a mammoth amount of storage, which will no doubt cost a bundle), the *Synclavier* will take the next evolutionary step beyond synthesisers and emulation machines. It will become a true digital recorder and more—it will in effect be a complete studio in one package, lacking only microphones and speakers. Wait for it. ■