The Kurzweil 150 Fourier Synthesizer

Sound Modeling Program™

By Hal Chamberlin

The Kurzweil 150 Fourier Synthesizer is a remarkable achievement in musical engineering with its full complement of authentic instrument sounds and uniquely complete MIDI implementation. Now with the Sound Modeling Program™, it becomes one of the most sophisticated programmable synthesizers as well. The Sound Modeling Program allows you to not only edit the preset sounds of the 150FS, but also create new sounds never heard before by manipulating envelope shapes and frequency spectra on the screen of the Apple IIe.

Of course, drawing waveforms on a computer screen is nothing new. Several samplers, including the Kurzweil 250, are supported by programs for doing just that. A few of these even allow additive synthesis to be simulated in software, and the result downloaded to the sampler to be heard.

The Kurzweil 150 Fourier Synthesizer, however, is a true additive synthesis machine. Sounds may not only be designed and edited in the frequency domain, they are really synthesized in that way. This allows very rapid program operation and produces sound files that are compact and load quickly. It also means there are none of the typical sampler problems of finding loop points and hearing noticeable repetition during long sustains.

HOW THE KURZWEIL 150FS PRODUCES SOUND

In the 150FS is a "pool" of 240 independent oscillators called partials, each of which may be a sine wave or any of several types of noise. When a note is played, the necessary number of partials is assigned from the pool. Each oscillator's output is then given its own arbitrarily-shaped amplitude envelope, called a contour, which comes from the sound model memory. The contours themselves are represented as a series of line segments where each breakpoint has a time (to millisecond resolution) and a rate or amplitude slope in decibels per second.

An instrument voice is typically made of many models, each of which covers a fairly narrow pitch range. Each model in turn has its own set of contours and a multi-level attack function which specifies how the dynamic spectrum defined by the contours is to be modified for different MIDI key velocities. Additionally, there are parameters in the models which specify how the contours are modified for long sustains and short staccatos.

WHAT THE SOUND MODELING PROGRAM DOES

The **Sound Modeling Program** primarily provides a graphic method for constructing and manipulating these data structures. In addition, it translates between a data representation that you can readily understand and the internal format used by the 150FS. For example, you display and manipulate contours in familiar time-amplitude form while they are communicated to the 150FS in time increment and slope form.

To the 150FS, a sound model is really a 3-dimensional surface combining time, frequency (partial number), and amplitude. For detailed editing on a CRT, only two of these dimensions can be dealt with at once. Of course, the Sound Modeling Program allows you to draw and edit envelopes of the partials individually, which means that you are working in the time-amplitude domain. But it also allows you to work in the amplitude-frequency domain where you specify the sound's spectrum at points in time, and the Sound Modeling Program will smoothly interpolate spectral changes between these points.

The Kurzweil 150FS Sound Modeling Program also provides carefully designed interactive screens for setting the partial frequencies (2 methods), attack function, and combining several models into a single, wide-range musical voice. Existing models, such as the standard factory sounds, can be read in from the 150FS and displayed or modified. In fact, anything that the Kurzweil 150FS can do can be specified using the Sound Modeling Program.

Normally, a Kurzweil MIDIBOARD® or other keyboard with MIDI output would be wired in for testing and playing the new sounds. The Sound Modeling Program, however, provides an "audition" function which uses the Apple keyboard for such testing. A unique "4X Slow" function temporarily slows the sound by a factor of 4 without affecting pitch for detailed aural study.

WHAT YOU DO

Configured for the Apple IIe, the Sound Modeling Program uses the digit keys as 10 "function keys" and the arrow keys to move a graphic crosshair block selection cursor around the screen. The current function name and what each digit key does in that function is always clearly labeled on the screen. Thus program operation consists mostly of pressing function (digit) keys and arrow keys, and occasionally entering names and numbers. When errors occur, informative messages are overlayed in a temporary message window. Even the storage, recall, and manipulation of sound model disk files is under function key control, thus minimizing your direct contact with the Apple][e operating system.

The Sound Modeling Program itself is written entirely in 6502 assembly language complete with its own high-speed graphic and compact text routines. Most functions execute in less than a second, while function key menu changes are practically instantaneous. Since additive synthesis sound models are so compact (typically under 1K each), they load into the 150FS very quickly for evaluation.

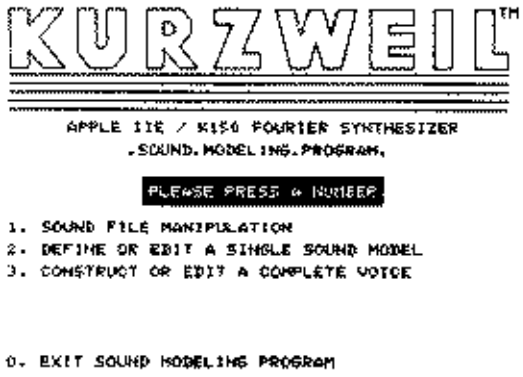

After loading, the Sound Modeling Program presents a menu of its major functions. Free movement among major functions is provided at this base level.

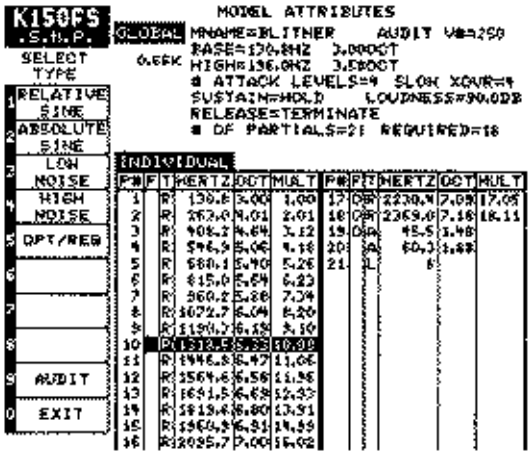

A complete listing of the global model attributes and the individual partial parameters is easily obtained. Here the user has selected the change partial type function. Cursor keys select the highlighted partial and digit keys initiate action.

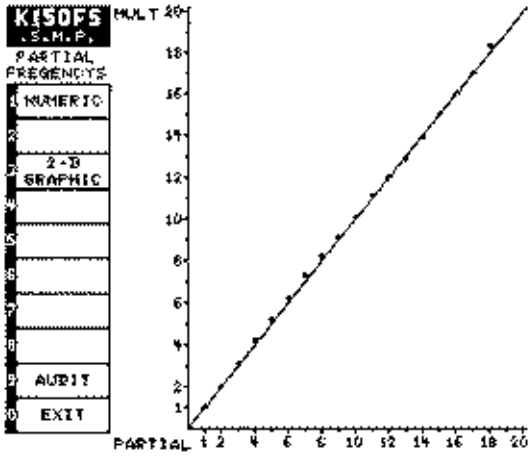

Partial frequencies need not be harmonically related. The 2-D graphic display clearly shows deviations from perfect harmonic frequencies. 'Selected' partials are indicated by solid dots.

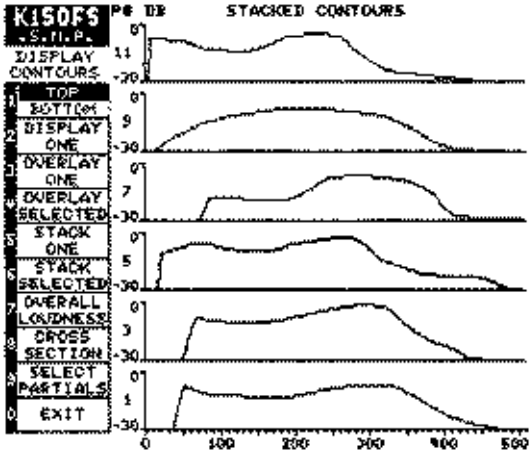

One method of working with dynamic spectral changes is to display the contours of selected partials as a 'stack' of plots. Detailed editing is performed on one such plot expanded to cover half the screen.

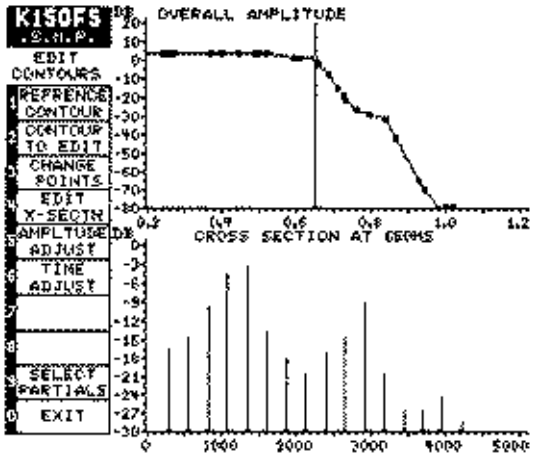

A cross section plot shows the sound's spectrum at any point in time. Note that time has scrolled forward to show the tail end of the sound's envelope. Breakpoints of all of the partials combined are shown on the overall amplitude plot.

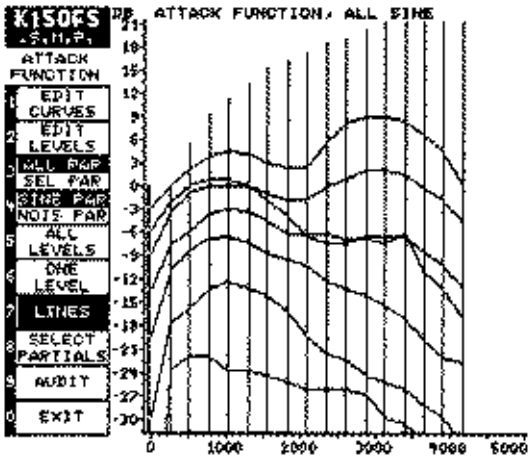

The attack function shows how the spectrum is modified for different loudness levels. The wide bar represents the overall loudness while the remaining vertical lines represent the amplitudes of selected partials. Connecting lines show the spectrum for each of the ''attack levels.''

The Kurzweil 150 Fourier Synthesizer

Sound Modeling Program

VERSION 1.0 SPECIFICATIONS

COMPUTER:	Apple IIe with DOS 3.3 and one disk drive; 80-column card not needed
ADDITIONAL HARDWARE:	Passport Designs MIDI interface card in slot 2
DISPLAY:	Apple II high resolution graphics (280 x 192), monochrome
NUMBER OF PARTIALS:	1 to 64 partials
MODEL SIZE:	128 bytes minimum, 7K bytes maximum (over 2000 breakpoints)
SOUND RAM:	64K bytes, non-volatile
MODELS PER VOICE:	1 - 127; typically 2 - 4 per octave
ATTACK LEVELS:	1 - 127; typically 4 - 10
LONGEST SOUNDS:	65.5 seconds; indefinite with looping
TIME RESOLUTION:	1 millisecond
TIME SCALES:	50, 100, 200, 500 ms, 1, 2, 5, 10 s full screen with scrolling
AMPLITUDE RESOLUTION:	0.375 dB at breakpoints, interpolated to 0.0015 dB
AMPLITUDE SCALES:	20, 30, 40, 50, 60, 80 dB full screen
FREQUENCY RESOLUTION:	0.298 Hz
FREQUENCY SCALES:	600, 1K, 1.6K, 2K, 3K, 4K, 5K, 6K, 8 kHz full screen
DISPLAY SCREENS:	Title, partial parameters (global and individual), partial frequencies (tabular, 2-D graphic), contours (individual partial, 2-17 stacked, line plots), cross-section (static and dynamic), attack function (with or without lines), release rates, audition, sound file listing.

KURZWEIL MUSIC SYSTEMS, INC.
Waltham, MA
April 1987