6.5 Synthesising new voices: the FFT page

The FFT (Fast Fourier Transform) page provides facilities for harmonic synthesis and analysis. To enter the FFT page,

FFT<CR>

FFT analysis is a method for alternating between two representations of a wave. The standard representation of a waveform is to plot amplitude against time (as in the WaveEdit Page). When a waveform is displayed in the FFT page however, it is expressed in terms of amplitude and phase of a harmonic series against time.

The harmonic series is based on a fundamental frequency equal to the sample rate decided by the analysed segment length (e.g 44100Hz / 2048 = 22Hz). to achieve a more useful display for pitched sounds, it is suggested that the pitch of the sound is a multiple of the fundamental frequency and the segment length is a power of 2. This can be achieved by using the pitch analysis and sample rate convert commands found on the WaveEdit page.

The Fourier transform can be viewed in two orientations:

Segment mode - Amplitude (or pitch) versus harmonic number With frame number changing with 3D depth.

Profile mode - Amplitude (or pitch) versus frame number changing with 3D depth.

Frame number is equal to segment number multiplied by the overlap ratio (see See Segment Length 6.5.3.2.). You can regard either parameter as being equivalent to time. While in segment mode, the original waveform can be also displayed as a function of time (in a similar manner to the WaveEdit page).

To modify a Subvoice on the FFT page is first necessary to analyse the Subvoice to calculate the Fourier transform of the Subvoice. You can then edit individual harmonics in the transform using the Draw mode displays and then create a new waveform using the Synthesis command. Once a Subvoice has been analysed, further changes to the transform require only the drawing and synthesis steps.

6.5.1. Definitions: Segment, Harmonic Frame, Harmonic Profile

6.5.1.1. Segment

A segment on the FFT page refers to a division of the waveform for analysis and synthesis purposes. It does not necessarily relate to the pitch of the waveform, as does segment length after pitch analysis on the sample page. (See P6.4.6). Segments on the FFT page may vary in length. They must however be powers of 2, between 32 and 4096. Existing waveforms may be modified to conform to these segment lengths using the Sample Rate convert command (P7.10). This is not necessary for analysis purposes, however. Segments are numbered from zero upward.

6.5.1.2. Harmonic Frame

The waveform is divided for analysis and synthesis purposes into frames. The length of each frame is equal to the segment length of the waveform. The first frame (frame zero), starts at the beginning of the waveform. Each subsequent frame is offset from the previous frame by a distance equal to the segment length divided by a parameter called the overlap.

$f_n = w + n*s/o$

where

f_n is the start of frame number n
w is the start of the waveform
n is the number of the frame
s is the FFT page segment length
o is the overlap

There are o times as many frames as segments. Thus for an overlap of 4, there are four times as many frames as FFT segments, and the position of frame 4 corresponds to the position of the segment 1.

The Harmonic Frame field displays the number of the currently displayed frame in Segment display mode.

6.5.1.3. Harmonic Profile

The Harmonic Profile is a graph of the amplitude of a given harmonic versus segment number. Segment number corresponds to time, so it is a graph of how the amplitude of the harmonic varies throughout the duration of the sound.

The Harmonic Profile field displays the number of the currently displayed Harmonic when in Profile display mode.

6.5.2. Analysing and Synthesising a Waveform

As each frame is successively analysed, the resultant frames are then crossfaded across the frame boundaries to create a smooth waveform. Consequently, the higher the overlap the smoother the resultant waveform. The other consequence of this smoothness is that the analysis and synthesis takes longer. Synthesis is the analysis process in reverse.

The FFT page requires an additional segment at the beginning of the waveform to permit it to overlap the first segment. This segment is automatically appended when the waveform appears in the FFT page.

6.5.2.1. Analysing a Waveform

The analyse command calculates the harmonic amplitude and phase representation of the sound from its waveform:

<F9>
or
AN<CR>
or

AN[HIT]

All segments must be analysed before the harmonic profile can be displayed.

6.5.2.2. Synthesising a waveform

This calculates the waveform representation from its harmonic amplitude equivalent:

SY<CR>
or
<F10>
or
SY[HIT]

When first used, the synthesis command synthesizes all the segments of a sound and then stores the harmonic amplitude information. Thus, for each segment of the sound the Series III stores the waveform, the harmonic amplitude and the phase representation.

6.5.3. Editing Synthesis and Analysis Parameters

The parameters window provides a facility for modifying parameters connected with the waveform size, segment size and analysis/synthesis modes. The parameters window is toggled on and off by the parameters command:

PA<CR>
or
<F1>
or
PA[HIT]

Changes to the parameters in the window have no effect until the parameters window is toggled Off by re-issuing the **PA**<CR> command. Once it is quit, the Series III attempts to make modifications according to the parameters set up. If it cannot do so, it displays an error message.

You must exit the Parameters window before attempting to display the Edit window (P6.5.6).

6.5.3.1. Number Segments :

Altering this field changes the size of the waveform (waveform length) by trimming or expanding it. The default value is the waveform size, measured in FFT segments. Changes to the waveform length affects the waveform as seen in other pages in the Series III. Similarly, moving the voice flags in the WaveEdit page will affect the waveform as seen on the FFT page. This is because the Trim function is automatically performed on the waveform before entering the FFT page.

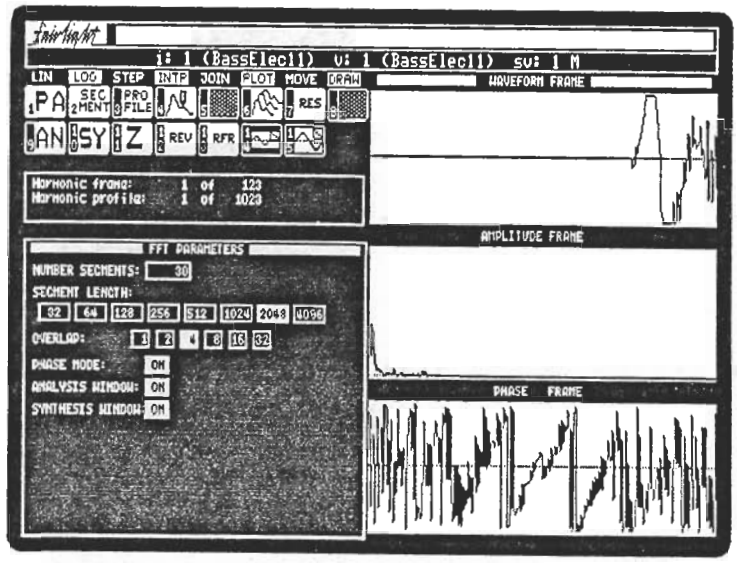

The FFT page showing the Parameters window after a Subvoice has been analysed

6.5.3.2. Segment Length

This field displays the allowed segment lengths for FFT page analysis and synthesis. The currently selected segment length is highlighted. To select a segment length, move the cursor to the desired field and use the <ADD> key to select that segment length, the <SUB> key to deselect it, or the <SET> key to toggle it on or off.

Recommended value: Default 2048. This gives a band width of 22.5 Hz to 22.5 kHz at a sample rate of 44.1kHz.

Note: Increasing the segment length reduces the time resolution of the FFT (number of frames), unless the overlap ratio is increased to compensate.

6.5.3.3. Overlap:

The Overlap field gives the number of FFT frames for each time segment when in Phase mode. In this mode, the FFT may be computed for more frames than there are time segments in the waveform, since FFT frames overlap each other (e.g by 3/4 of their length if overlap is 4).

Overlapping increases the time resolution of the FFT and avoids introducing clicks in the reconstructed waveform. Phase mode and overlapping should be left ON unless you want a special effect.

Recommended value: Default 4

6.5.3.4. Phase Mode:

Default: ON
Recommended: ON

6.5.3.5. Analysis Window

The Analysis window is not a display feature, but a technical term in FFT. It is an operation performed on the waveform which reduces numerical artifacts. This should be left ON except for special effects.

6.5.3.6. Synthesis Window

The Synthesis window is the inverse FFT equivalent of the Analysis window and should be left ON.

6.5.4. Segment Display

The segment display is the default display mode of the current waveform. Segment display shows amplitude and harmonic segments. To select the Segment display when in Profile mode:

SG<CR>
or
<F2>
or
SEGMENT[HIT]

The displayed waveform is determined by assignment to the Harmonic Frame field below the icon bank. The meaning of harmonic frame is explained above.

Harmonic Frame: *framenumber*<SET>

6.5.5. Profile Display

The Profile display is similar to the Segment display, but shows only the amplitude and phase for a selected harmonic. The current harmonic is indicated in the Current Profile field.

To select the Profile display:

PR<CR>
or
<F3>
or
PROFILE[HIT]

To select the Current Harmonic Profile:

Harmonic Profile: *Harmonicnumber*<SET>

6.5.6. Displaying and Editing

To edit a selected harmonic curve or segment, toggle the Edit window ON:

DRAW<CR>
or
<F4>
or
Edit[HIT]

The display and edit window can be manipulated by assignment or by Gpad. The value of a selected point in the display is displayed in the **data:** field, at the top of the edit display. This is the value corresponding to the position of the marker cursor, a vertical line which defaults to the far left of the window. The position of the marker is altered by [HIT]ting the edit display, or by assigning a position value to the **marker** field at the top of the display.

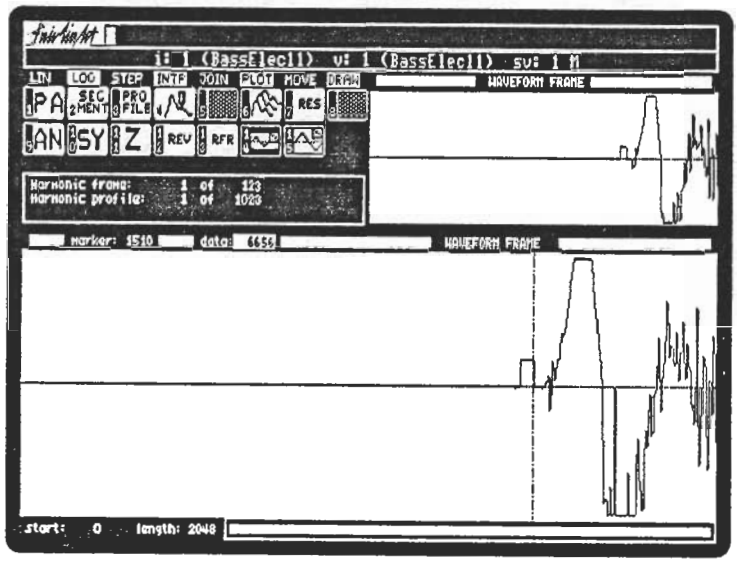

The Edit window showing a harmonic being altered with the Gpen

6.5.6.1. Selecting Curve to Edit or Display

The current curve displayed in the Edit window can be toggled between Phase and Amplitude mode with the following command:

<F15>
or
Big Sine Curve[HIT]

You can compare the two graphs by using the small ancillary display at the top right of the display:

<F14>
or
Small Sine Curve[HIT]

6.5.6.2. Drawing a Curve

Curves can be drawn or modified only in the large edit window not in the smaller display window. To draw a curve, select the desired drawing mode (see the Mode Menu section below), and draw by [HIT]ting the Display/Edit window. Alternatively a curve may be created by alphanumeric assignment using the marker, data, start and end fields around the border of the display.

start: The start point, either waveform position, (waveform segment) harmonic number (harmonic amplitude), or segment number (harmonic profile).

length: The length of the display in the appropriate units.

marker: The position in the display of the vertical cursor. This marker marks the position in the display corresponding to the value display.

To control the size and position of the display, assign to the **start** and **length** fields.

To assign a value to a point on the graph, assign the desired x (horizontal) value to the marker field, then assign the desired y (vertical) value to the data field at the right of the window base.

marker: *x*<SET>

data: *y*<SET>

Note: When modifying the time source of a harmonic in profile mode (particularly if reducing its amplitude), you should remember to modify the adjacent harmonics in a similar manner. If this is not done, then the harmonic may be frequency shifted rather than reduced in amplitude.

6.5.6.3. Mode Menu

The mode menu appears above the icon bank at the top left of the screen when the edit window is selected. The options are selected by [HIT]ting the option name, or by using <SHIFT> function keys. The options appear above the function key icon corresponding to their shifted function key.

Linear drawing mode:

LIN<CR>
or
<SHIFT><F1>
or
LIN[HIT]

Logarithmic drawing mode:

LOG<CR>
or
<SHIFT><F2>
or
LOG[HIT]

The following drawing facilities can also be used.

In step mode, the display shows a horizontal step corresponding to the value at each of these points:

STEP<CR>
or
<SHIFT><F3>
or
STEP[HIT]

In Interpolate mode, the curve is smoothly interpolated between these points:

INTP<CR>
or
<SHIFT><F4>
or
INTP[HIT]

In Join mode, the curve is joined between the point just plotted and the previous point plotted:

JOIN<CR>
or
<SHIFT><F5>
or
JOIN[HIT]

In Plot mode, no interpolation takes place and the new point is plotted without alteration to other points:

PLOT<CR>
or
<SHIFT><F6>
or
PLOT[HIT]

In Move mode, a [HIT] on the curve moves the marker to that point without plotting the point.

MOVE<CR>
or
<SHIFT><F7>
or
MOVE[HIT]

In Draw mode a [HIT] on the curve plots that point.

DRAWM<CR>
or
<SHIFT><F8>
or
DRAW[HIT]

6.5.6.4. Scroll bar display control

The scroll area appears at the base of the edit window. Inside the scroll area is a highlighted scroll bar. The position of the scroll bar in the scroll area corresponds to the extent of waveform being viewed. For example, if the scroll bar fills only the first 1/3 of the scroll area, then only the first 3rd of the curve is in view in the display/edit window.

Touching the scroll area with the Gpad cursor brings a virtual bar into being. This virtual bar may be manipulated into the desired position. A [HIT] will then reposition the scroll bar at the virtual bar's position. If the Gpad cursor is moved out of the scroll area without a [HIT], the scroll bar will not change position.

If the Gpad cursor touches the scroll area outside the scroll bar, the virtual bar will extend to the Gpad cursor position, leaving its other end fixed. Moving the Gpad cursor then moves only the mobile end of the bar.

If the Gpad cursor touches the interior of the scroll bar, the whole bar will move with the Gpad cursor.

6.5.6.5. Zeroing a segment or profile

To zero the displayed curve:

Zero<CR>
or
<F11>
or
Z[HIT]

If not in edit mode, the curve zeroed will be the one with the highlit header line. See P6.5.6.1 to select this curve.

6.5.6.6. Copying a segment or profile to other segments or profiles

[Not available this release]

6.5.6.7. Reset the display

<F7>
or
RES[HIT]

This resets the display to original segment start and length settings for the horizontal and depth axis.

6.5.6.8. Refresh the display

This will alter the display parameters:

<F13>
or
RFR[HIT]

6.5.7. 3D display mode

The 3 dimensional displays a waveform either in Profile or Segment mode. Segment mode works like the Waveform Edit page 3 D display, displaying consecutive segments of the waveform behind on another. The Profile mode displays the harmonics of the waveform one after another.

By adjusting the parameters in the box above the 3D display, it is possible to view a selected range of harmonics or segments. The table below describes the parameters available.

To enter 3D display:

TD<CR>
or
<F6>
or
stacked waves [HIT]

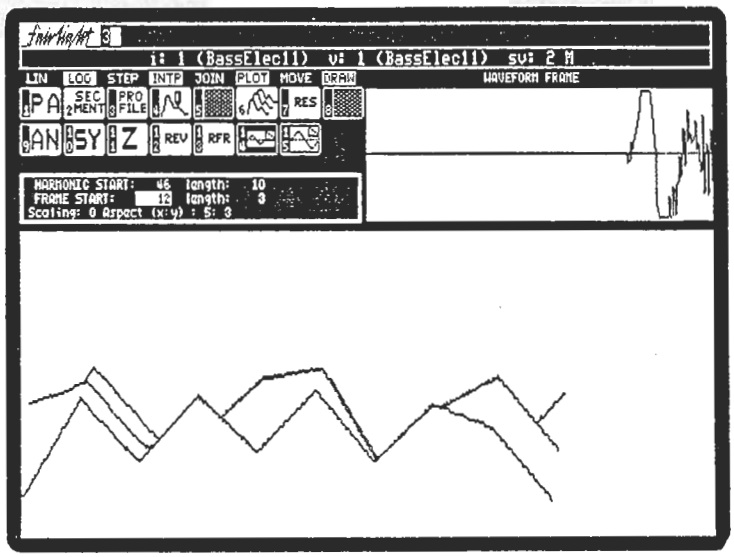

3D window showing only 3 segments of Subvoice. The length of the waveform has been reduced to 10 to help define the shape of the segment.

In Segment mode:

Field name	Axis	Comment
Harmonic Start:	Horizontal	Harmonic number (Freq. as a multiple of FFT fundamental frequency).
Frame Start:	Depth	Frame number (same as time).
length:	The upper length field controls the 3D depth. The lower length field controls time.	

In Profile mode:

Field name	Axis	Comments
Frame Start:	Horizontal	Frame number (same as time).
Data Start:	Depth	Shows the harmonics beginning from the number in this field.
length:	The upper length field controls the 3D depth. The lower length field controls time.	

6.5.7.1 Scaling :

The Scaling parameter can be used to increase the vertical resolution of the display by a factor of 2. Display clipping may occur if the factor is too large.

6.5.7.2. Aspect Ratio:

The aspect ratio displays the number of pixels (dots) the display of each subsequent segment sector is displaced to the right (x) for a corresponding upwards displacement (y). An aspect of 0:3 gives a vertical display stack. An aspect of 3:3 gives an oblique angle display of about 45 degrees.

Note: It is suggested that you step through the profile display at a low resolution (less the 30 depth steps at a aspect ratio of 4:2). This helps view as many of the peaks as possible which may be other wise missed if the resolution is too high.

6.5.8. Creating a New Subvoice

To create a new Subvoice in the Current Voice:

SVC<CR>

The new Subvoice becomes the current Subvoice. It contains the default sine waveform, which cannot be edited as it is the common property of all new default Subvoices. This Subvoice must be expanded before synthesized segments can be added to it.

6.5.9. Expanding a Waveform

To increase the size of the Subvoice waveform, you use the SVX command.

SVX<CR>
or
SVX><CR>

adds 1k of waveform to the tail of the waveform.

SVX*n*<CR>
or
SVX*n*><CR>

adds *n*k to the tail of the waveform.

SVX<<CR>

adds 1k to the head of the waveform.

SVX*n*<<CR>

adds *n*k to the head of the waveform.

e.g. To create a new Subvoice with room for 10 synthesized segments:

SVC<CR>
SVX10<CR>