

# Harmonic visualizer

## Quick start guide

### **What is this document about?**

This document is a detailed introduction to harmonic visualizer (HV) for editing audio using harmonic sinusoidal modelling techniques.

### **Who is this document for?**

This document is for users of HV with general knowledge about audio waveforms, spectrograms and harmonic structures of pitched sounds.

### **What is needed to run the program?**

PC with Window operating system and basic audio I/O devices.

### **How do I install the program?**

The program is shipped as a stand-alone executable and does not require installation.

### **How do I uninstall the program?**

The program does not use the system registry. The program folder can be deleted directly.

### **What audio formats are supported?**

Single- or dual-channel 8-, 16- or 24-bit linear PCM waveform contained in .wav files.

# Table of contents

1. Introduction .....	4
2. Knowing the main window .....	5
2.1 The toolbar .....	6
2.2 The WaveView .....	7
2.2.1 The display panes .....	8
2.2.2 The fixed tags .....	9
2.3 The navigator window.....	10
2.4 The setting panel.....	11
2.4.1 Display settings tab.....	11
2.4.2. Playback settings tab .....	12
2.4.3 Harmonic sinusoid tracking settings tab .....	13
2.5 The textual spectrogram display .....	14
3. Display options.....	15
3.1 Waveform / Spectrogram.....	15
3.2 Channel selection .....	16
3.3 Spectrogram display settings.....	17
3.4 Axes and units.....	17
3.5 Cursor texts and pane texts .....	18
3.6 Spectrogram parameters .....	18
4. Time and frequency selection .....	19
4.1 Select a duration .....	19
4.2 Select a band .....	19
4.3 Select an area .....	19
4.4 Resize/relocate a selection.....	20
4.5 Multiple selections .....	20
4.6 Time-frequency masking .....	20
5. Navigation .....	21
5.1 Select and zoom.....	21
5.2 Mouse wheeling.....	21
5.3 Scrollbar .....	22
5.4 Mouse dragging .....	22
5.5 Range tags .....	22
5.6 The navigator window.....	22
5.7 Restore.....	22
6. Playback .....	23
6.1 Channel selection .....	23
6.2 Playback filters .....	23
6.3 Looped playback.....	24
7. Recording .....	25
8. Working with harmonic sinusoids .....	26
8.1 Event viewer.....	27
8.2 Selecting harmonic sinusoids .....	28
8.2.1 Partial selector .....	28


8.2.2 Algorithmic parameters of harmonic sinusoid tracking .....	28
8.2.3 Manual termination and continuation .....	28
8.2.4 Modifying selected events .....	29
8.3 Editing harmonic sinusoids.....	30
8.3.1 Extracting/removing harmonic sinusoids .....	30
8.3.2 Editor panel.....	30
8.3.3 Amplification.....	31
8.3.4 Pitch shifting .....	31
8.3.5 Amplitude and frequency modulations.....	32
8.3.6 Frequency demodulation and remodulation .....	32
8.3.7 Saving the editing result.....	33
8.4 Vibrato analysis and synthesis demo .....	34
8.4.1 GUI .....	34
8.4.2 The pitch trajectory panes .....	35
8.4.3 The vibrato representation panel.....	36
8.4.4 Source-filter modelling settings .....	38
8.4.5 Other controls .....	38
8.5 Source-filter modelling demo .....	39
8.5.1 GUI .....	39
8.5.2 The pitch trajectory panes .....	40
8.5.3 The source-filter representation panel.....	41
8.5.4 Source-filter modelling settings .....	42
8.5.5 Display settings .....	42
8.5.6 Saving and using source-filter models .....	43
8.5.7 Other controls .....	43
References .....	44

# 1. Introduction

This document provides a quick start guide for using HV for editing harmonic sinusoidal components in audio. HV comes as a Windows .exe application and works with 16-bit PCM waveform files, mono or stereo. A mouse with roller wheel is recommended for smooth use. The current package ships with the following files:

hv.exe	main application executable
hv.ini	program settings, readable and editable
451spfvf-1.wav	example audio file
QuickStart.1107.pdf	this file


HV can work on no more than one audio file at each time. By default when the program starts it loads in the audio file that was open when the program was last closed. To load another file use File|Open from the main menu, or the Open File button leftmost on the tool bar.



When the program is first launched the example audio file is automatically loaded, as indicated in the settings file.

## 2. Knowing the main window

HV enters the main window upon start. The working area of the main window contains a main menu at the top, a toolbar just below it, a graphical waveform/spectrogram display panel (the *WaveView*, 1a) on top-left, a textual spectrogram display panel (1b) on bottom-left, a scrollbar below the WaveView, a navigation window (the *Navigator*, 2) on top-right, a setting panel on the right (3), and a status bar at the bottom.



This button is used to hide or show (2) and (3).

This button is used to hide or show (1b).

## 2.1 The toolbar

The toolbar currently contains 12 tool buttons and 3 combo-selection boxes as shown below.

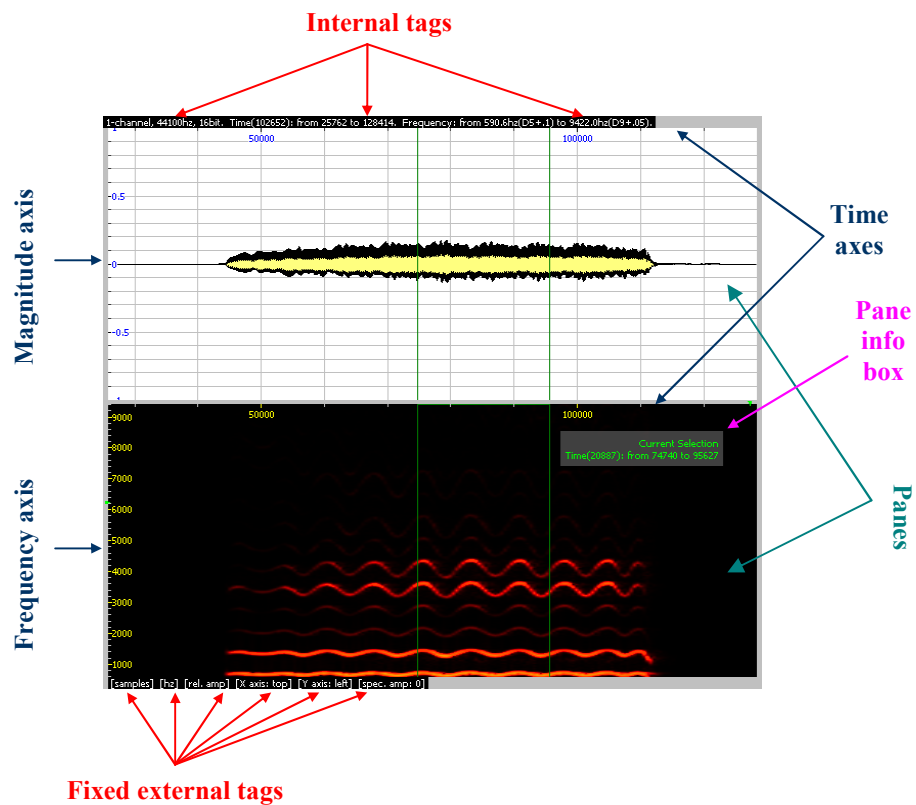


Table - ToolBar buttons

Name	Equiv. menu item	Description
File Open	File Open	Open a file
File Save As	File Save As	Save the current visible audio as a file
Exit	File Exit	Quit program
Record		Start or Stop recording
Play		Start or stop playback
W/S Swith		Toggle waveform / spectrogram views
Cursor Test		Show or hide info text at cursor
Pane Info		Show or hide info text inside main display area
Time Selection		Switch time selection on / off
Band Selection		Switch frequency selection on / off
Multiple Selection		Switch multiple selection on / off
Harmonic Selection		Set harmonic sinusoid select mode


## 2.2 The WaveView

The WaveView features one or more display panes, arranged according to display settings, along with a number of *tags*. The WaveView panes are areas for displaying waveforms or spectrograms, along with supporting elements such as axis, gridlines, cursors, and info boxes. *Tags* are visual elements attached onto the WaveView used for displaying quick information as well as accepting specific user inputs.




## 2.2.1 The display panes

The waveform pane displays the audio content as raw samples, in which the horizontal axis refers to time, and the vertical axis to the sample value (corresponding to microphone membrane displacement, sound pressure, etc.). The time axis is marked in seconds or samples. The magnitude axis is marked in absolute (e.g. -32768~32767 for 16bit waveform) or relative (i.e. -1~1 for any format) values. Additional information that does not require frequency alignment can be displayed in the waveform pane.



The spectrogram pane displays the audio content as an energy distribution in the time-frequency plane. The horizontal axis refers to time, and the vertical axis to frequency. The time axis is marked in seconds or samples. The frequency axis is marked in Herz or DFT bins. Additional information that requires frequency alignment can be displayed in the spectrogram pane.



See 3 “Display options” for how to show or hide the waveform or spectrogram panes.



## 2.2.2 The fixed tags

Tags are small labels tagged onto the WaveView to display quick information and to accept user inputs, usually regarding the objects whose quick information each tag is displaying. There are many types of tags used in HV. Nine of these tags are fixed tags and are always displayed.

At the top of the WaveView are three “internal” tags: a format tag on the left, time range tag in the middle, and frequency range tag on the right. The format tag displays the PCM format version of the current audio content. The time and frequency range tags display the time and frequency ranges currently shown in the panes. By double-clicking the time and frequency range tags one can type in new display ranges to be shown in the WaveView panes. See 3. “Display options” for more info.

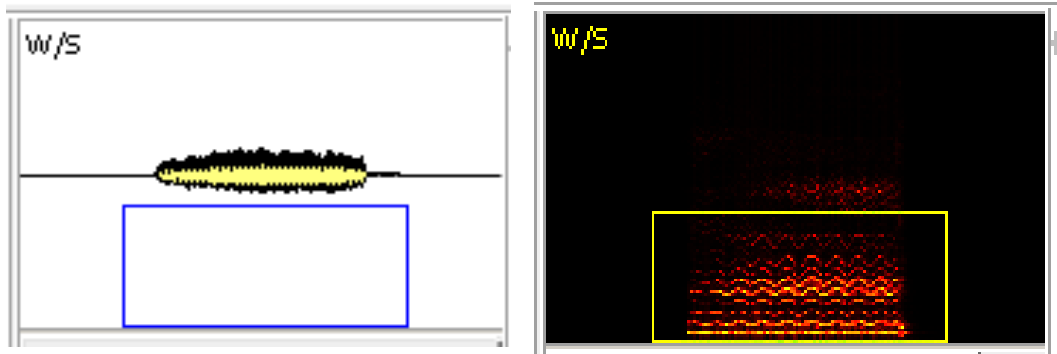
1-channel, 44100hz, 16bit, Time(102652): from 25762 to 128414, Frequency: from 590.6hz(D5+.1) to 9422.0hz(D9+.05).

At the bottom of the WaveView are six “external” fixed tags: (from left to right) the time unit tag, frequency unit tag, magnitude unit tag, X axis tag, Y axis tag and spectrogram brightness tag. The three unit tags are used to control what units (scales) are used for displaying and typing-in time, frequency and sample magnitude. The X and Y axis tags turn the axis display (and in the case of waveform view, the grid lines) on and off, as well as control the positional of the axes being displayed. The spectrogram brightness tag is used to control the brightness of spectrogram display. The first five tags accept single mouse click input; the spectrogram brightness tag accepts mouse wheeling input. See 3. “Display options” for more info.

[samples] [hz] [rel. amp] [X axis: top] [Y axis: left] [spec. amp: 0]

## 2.3 The navigator window

The navigator window displays a miniature image of the current audio file in waveform or spectrogram. The time and frequency range currently used for displaying the audio content in the main WavView is shown by a rectangular frame. This frame can be dragged or resized to change displaying ranges. See 4. “Navigation” for more info.

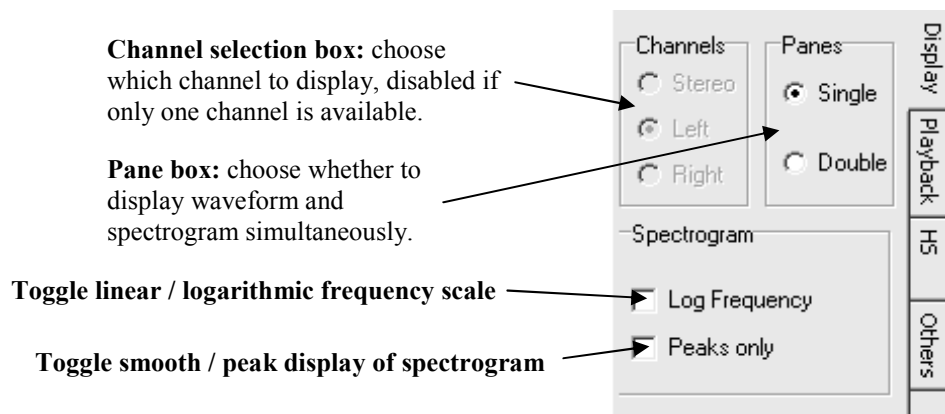


## 2.4 The setting panel

The setting panel contains 4 tab pages, controlling the program settings regarding display, audio playback, harmonic sinusoidal tracking, etc.

### 2.4.1 Display settings tab

The display page controls how waveforms or spectrograms are displayed in the WaveView panes.



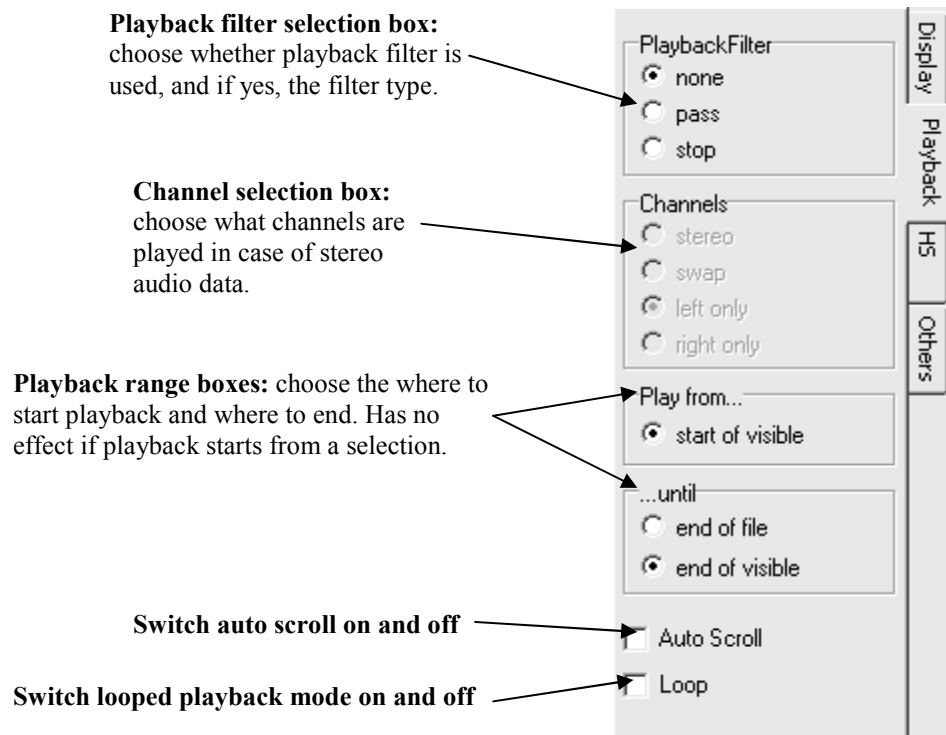
The channel selection box controls which channel (or both channels) of a stereo input file to display in the WaveView. The pane box controls whether to display both waveform and spectrogram or to display either of them.

The log frequency checkbox toggles the frequency axis in spectrograms between linear and logarithmic scales. The “peak only” checkbox toggles the spectrogram display between normal and peak-only mode, in which only local spectral peaks are drawn.

For more info see 3. “Display options”.

## 2.4.2. Playback settings tab

The playback setting page controls how the audio content is played. Aside from normal playing, the user has the choice of applying time-frequency masks during playback, remapping playback channels of a stereo audio, or play in loops.



Playback filters allow playing audio through time-frequency masks, i.e. only components inside or outside a T-F mask is played. The mask can be an interval, a pass band, a time-frequency box, or a combination of such areas in the T-F plane. The playback filter selection box control whether to use playback filters, and whether to play from inside or from outside the T-F mask.

Channel selection box controls which channel(s) of a stereo audio to play. In case of both channels are played, there is the optional of playing swapped channels.

The playback range boxes control the duration of audio to play, the options being from the beginning to the end of the displayed range, or from the beginning of the displayed range to the end of file.


The “auto scroll” checkbox controls whether the display is scrolled forward down time axis to follow the audio being played when playback goes out of the displayed time range.

The “loop” checkbox controls whether playback terminates at the end of the duration being played, or restart from the beginning into a loop.

See 6. “Playback” for more info.

## 2.4.3 Harmonic sinusoid tracking settings tab

The harmonic sinusoid analysis settings page (shortened as HS on the tab) contains a partial index selector and six algorithm parameter entries. The partial selector is used to supply the harmonic sinusoid tracker with its partial index input, as well as to display which partial is currently being worked on. The algorithmic parameters are used to control the behaviour of the harmonic sinusoidal tracking algorithm. For more info see 8. “Working with harmonic sinusoids”.



## 2.5 The textual spectrogram display

tick this to allow automatic scrolling only if Shift key is held down


floating-point number format


Amp. Spec.	Phase Spec.	Quick Peaks							<input checked="" type="checkbox"/> %3f	<input type="checkbox"/> Shift
	225	226	227	228	229	230	231	232		
108	1.138	0.381	1.301	0.553	0.436	0.365	0.127	0.244		
107	0.614	0.811	0.797	0.393	0.187	0.253	0.263	0.186		
106	1.184	0.941	0.607	0.401	0.168	0.217	0.091	0.194		
105	2.228	1.008	0.695	0.563	0.257	0.288	0.055	0.251		
104	1.096	0.866	1.303	0.824	0.133	0.288	0.126	0.321		
103	1.123	0.754	1.371	0.596	0.154	0.110	0.116	0.250		
102	1.141	1.443	0.911	0.309	0.167	0.164	0.155	0.494		
101	0.431	1.716	0.625	0.247	0.249	0.087	0.214	0.358		
100	1.454	0.708	1.192	0.563	0.184	0.107	0.684	0.296		
99	2.515	1.653	1.467	0.836	0.270	0.240	0.656	0.447		

The textual spectrogram display allows direct numerical observation of the amplitude and phase spectrograms in two tables. Double-clicking a table grid sends the mouse pointer to the position indicated by the (frame, bin) indices to that grid. A third table shows spectral peak measurements (frequency and amplitude) using the QIFFT method.


## 3. Display options

### 3.1 Waveform / Spectrogram

The Waveform/Spectrogram Switch button () switches between the waveform and spectrogram displays. Frequency-related operations are not available on waveform display.



Set Display panes setting to “Double” to display waveform and spectrogram at the same time.



## 3.2 Channel selection

Use the Channel selection box on the Display setting page to choose which (or both) channel to display if the audio format is stereo.

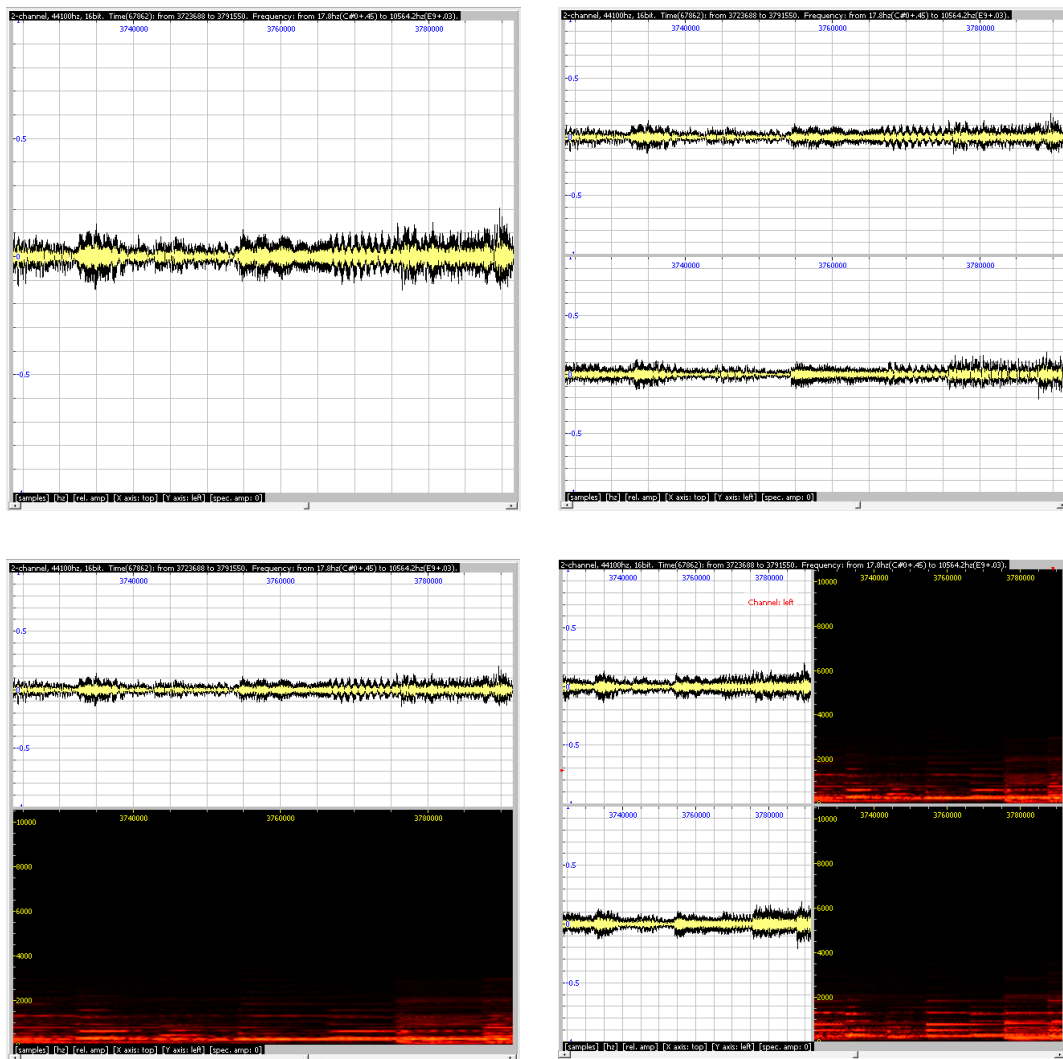


Figure: arrangement of displaying panes  
(mono - waveform display)(stereo - waveform display)  
(mono - double display)(stereo - double display)



### 3.3 Spectrogram display settings

The Window setting box on the Display setting page controls the type or size of the window function used for calculating the spectrogram. These settings are also used in other DSP modules including the sinusoid modeling.

The frequency scale can be linear or logarithmic as indicated by the “Log frequency” checkbox on the Display setting tab page. The peak spectrogram is displayed instead of plain spectrogram if “Peaks only” checkbox is checked.

The brightness of spectrogram is controlled by mouse-wheeling the spectrogram brightness tag, displayed with [spec.amp]. Double-click this tag to restore default value 0.

[samples] [hz] [rel. amp] [X axis: top] [Y axis: left] [spec. amp: 0]

Automatic spectrogram brightness is enabled/disabled by double-clicking the spectrogram brightness tag while holding down the Ctrl key. In this mode the brightest point in the display is held at a constant brightness.

[samples] [hz] [rel. amp] [X axis: top] [Y axis: left] [spec. amp (auto): -0.0144]

### 3.4 Axes and units

The X and Y axes are located on the top and the left by default. They can be relocated to the bottom / the right or turned off by clicking the “X axis” and “Y axis” tags, found in the figure above.


[samples] [hz] [rel. amp] [X axis: top] [Y axis: left] [spec. amp: 0]

The first three tags in the same row are unit tags, by clicking which the user can change the units used for displaying time, frequency and magnitude. The unit settings also affect input formats of these variables throughout the program.


[samples] [hz] [rel. amp] [X axis: top] [Y axis: left] [spec. amp: 0]

### 3.5 Cursor texts and pane texts

When moving the pointer the current position of time/frequency/magnitude at the pointer is marked by triangular floaters against the x and y axes. Cursor texts are displayed beside the floaters to show the reading of cursor positions. Use the Cursor Text button (0.16s) to switch cursor text on or off.




Pane info texts are displayed in the main display area to provide basic information of the audio content currently being displayed in each pane, the selection currently focused on, or sinusoidal atoms currently being pointed at (see 8. “Working with harmonic sinusoids”). Use the Pane Info button (sam 224) to switch pane info text on or off.





### 3.6 Spectrogram parameters




Use the Window Type and Window Size selection boxes to change the analysis window used for computing the spectrogram. Notice that this window size is also used in the harmonic sinusoidal analysis process.





## 4. Time and frequency selection

The program features two selection modes: a normal selection mode in which the user can mark out an interval, frequency band, or time-frequency box with standard mouse-dragging, and a harmonic sinusoid selection mode in which the user can select a harmonic event. To work with in normal selection mode, make sure the Harmonic Selection button () is not pressed down.





In normal selection mode one may make a conventional selection in time (select a duration), frequency (a band), or both time and frequency (band-limited duration, or area). These are controlled by the Time Selection () and Frequency Selection () buttons. However, no matter if the user is selecting in time or frequency or both, the actual selection is a bounded both in time and in frequency. One can also make multiple selections using the Multiple Selection () button.



### 4.1 Select a duration



Make sure the  button is down and  button is up. Drag mouse pointer to select. One may selection a duration both in waveform and in spectrogram displays. The selection is bounded in frequency by the current frequency range of the spectrogram display in both cases, even if the spectrogram display is not visible.

### 4.2 Select a band

Make sure the  button is down and  button is up. Drag mouse pointer in the spectrogram display to select. The selection is bounded in time by the currently displayed time range.

### 4.3 Select an area




Made sure both  and  buttons are down. Drag mouse pointer to select.

Tip: You need to hold the Shift key down when pressing either  or  button without popping up the other.

## 4.4 Resize/relocate a selection

A selection can be resized by holding the Ctrl key and drag its boundary or corner. A selection can be relocated by holding the Ctrl key and drag its body.

## 4.5 Multiple selections

The  button controls whether multiple selections are allowed in the normal selection mode. When  button is up, all selections are cancelled whenever left mouse button is pressed, which ensures no more than one selection is present. Press  button down to make new selections without cancelling the existing ones.

When multiple selections are present, there is one “focused” selection marked by solid boundaries while the others have dashed boundaries. The focus can be switched by clicking inside the selection to be focused at, or by PageUp and PageDown keys.

The focused selection can be removed by Delete key. To clear all selections, press Esc.

## 4.6 Time-frequency masking

Single or multiple time-frequency selections constitute a binary mask in the time-frequency plane. When a time-frequency mask is present, one can extract or remove the contents inside the mask using Extract or Cut commands from the popup menu of the main spectrogram display. These operations do not affect audio data outside the currently displayed duration.

## 5. Navigation

*Navigation* refers to changing the time and frequency ranges that is currently being displayed in the WaveView. HV provides various ways to navigate the audio content in the time-frequency plane.

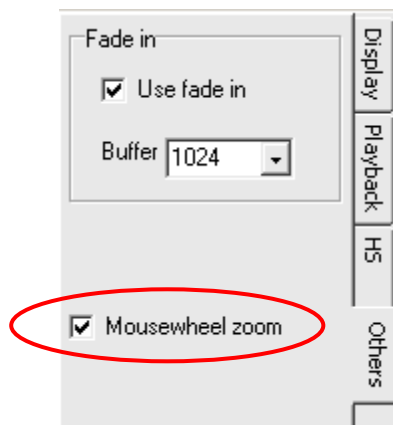
### 5.1 Select and zoom

The quickest way to zoom-in to a specific part of the T-F plane is to make a selection (duration, band or area) and use “Zoom to selection” from the popup menu. In case of multi-selection, the focused selection with solid frame is used.

By default a selection made before a zoom-to-selection operation is considered only for zoom-in purpose and is therefore removed after zooming. To keep this selection after zooming, hold the Shift key down when clicking “Zoom to selection” from menu.

### 5.2 Mouse wheeling

Mousewheel zooming is enabled (default) by checking the “Mousewheel zoom” checkbox on the Other settings tab page.



Use the mouse wheel alone to zoom in or out in time. In waveform view use the mouse wheel with the Shift key (or Ctrl key or any mouse key) held down to zoom in or out in magnitude. In spectrogram view use the mouse wheel with the Shift key (or Ctrl key or any mouse key) held down to zoom in or out in frequency. In case of time/frequency zoom the position of the mouse pointer is used as the zooming centre; in case of magnitude zoom the zooming centre is always 0.

## 5.3 Scrollbar

A scrollbar is visible when the current display does not cover the whole time range of the audio. Drag scrollbar to shift the time range.

## 5.4 Mouse dragging

Hold the right mouse key down and drag inside the display pane to shift the time/frequency range being displayed.

## 5.5 Range tags

The time and frequency range tags are found at the top of the WaveView, above the panes. Double-click them to type in the new time or frequency range to display.

## 5.6 The navigator window

A miniature image of the current audio file is shown in the navigator window, with the current displayed range marked out by a rectangular frame. Use the W/S button at the upper-left corner of the navigator window to switch between waveform and spectrogram views.


The current display range can be moved about by dragging the body of rectangular frame inside the navigator window. The current display range can be expanded or shrunk by holding down Ctrl key and dragging the boundaries of the rectangular frame.

Tip: If the frame is too small to get at by moving the pointer, a right-click inside the navigator window will locate the pointer to the frame centre.

## 5.7 Restore

Use Restore from the popup menu to zoom to the whole time or frequency range.

## 6. Playback

Use the Play button () or Play from the popup menu to play the current audio.



If the playback is started using the popup menu which has been popped-up from inside the focused selection, then the audio is played from the start of this selection to the end of it. Otherwise the duration of playback is set as specified on the Playback setting tab page. A playback cursor is displayed during playback to mark the current playback position. If this cursor goes out of the visible range and “Auto Scroll” is selected, the program “turns the page” to follow the playback position.

### 6.1 Channel selection

If the current audio file is stereo one may choose to play either or both channels using the Channels box on the Playback setting page.

### 6.2 Playback filters

One may choose to play the audio content in its original form, or through a time-frequency mask. A time-frequency mask is specified by one or multiple selections. To apply the time-

frequency mask, use the Playback Filter selection box on the toolbar. Choose “pass” to play contents covered under the mask, “stop” to play contents outside the coverage of the mask.




## 6.3 Looped playback


When looped playback is enabled (see 2.4.2 for how to enable/disable looped playback), the current played content is looped.

If the playback is started using the popup menu that has been popped-up from inside the currently focused T-F selection, the looped interval goes from the start of this selection to the end of it. Otherwise the looped interval is as specified on the Playback setting tab page.




## 7. Recording

Use the Record button () to start recording from WAVE\_MAPPER device. A recording box is displayed during recording, which also has a Record button. The two Record buttons work interchangeably as one.




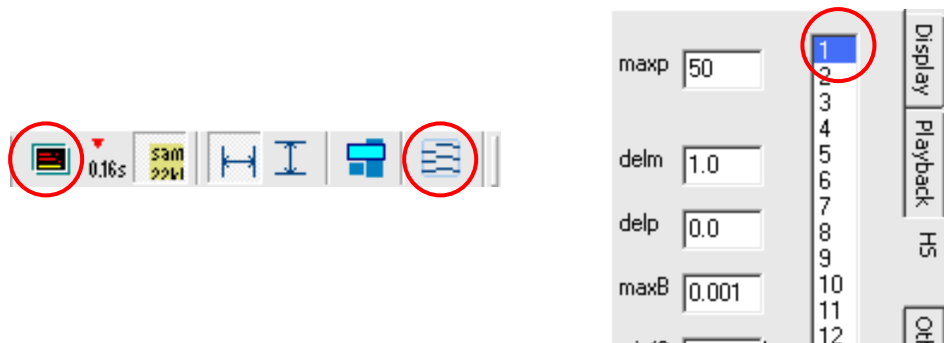
Recording is terminated using either Record button. User has the choice to discard the recorded content or send it to the main window as the current working file.




## 8. Working with harmonic sinusoids



The unique feature of HV is that it allows tracking and editing harmonic sinusoids independently of the concurrent audio contents. Harmonic sinusoids are sinusoidal representations of the deterministic components of pitched audio events. With HV the user can track, extract, remove, edit, as well as performing more specific vibrato and source-filter analyses of, the harmonic sinusoids.

Press the Harmonic Selection button  down to enter harmonic sinusoid selection mode. It is recommended to work only with the spectrogram view using a well-chosen window size. To get started, use the example audio, set the window type to Hann and size to 1024, make sure the “1” is selected in partial selector on the HS setting tab page, and click the lowest partial in the spectrogram.



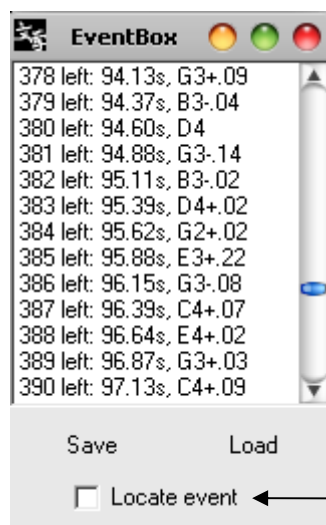
The harmonic sinusoid selected is displayed as atom tags, each of which shown by a “+”, with the exception of the atom at the user-specified point being shown by a “+” within a square. Moving the mouse pointer onto an atom tag highlights the whole selection. Hold Ctrl key down to highlight only one partial, or Shift key down to highlight only one frame.




Tip: Mouse clicks or dragging operations in harmonic selection mode can easily cancel or modify the current harmonic selection. If this is unwanted it is advisable to keep the  button up once harmonic selection is completed. Alternatively, the user can enter harmonic selection mode by holding down the S key instead of the  button.

## 8.1 Event viewer

The Event viewer is automatically displayed when a harmonic sinusoid is selected, or manually brought out using View|Events from the main menu. The Event viewer contains a list of selected harmonic sinusoids, labelled by the channel (left or right) and the time and frequency of the fundamental atom at the first frame. Select an item from the list to indicate which harmonic sinusoid is currently “active”. Press Delete key to remove the currently active item. Click on the white space or the box body to de-select all items before selecting a new harmonic sinusoid. This list can be saved with the Save button for future use.



Use this to follow selected event



The Event viewer has its own popup menu which is displayed by right-clicking an item in the event list. This menu contains entries to some basic editing operations as well as vibrato analysis/synthesis and source-filter modelling demos.

## 8.2 Selecting harmonic sinusoids

Harmonic sinusoids are selected by clicking the spectrogram on a sinusoidal partial that belongs to a harmonic event. This click supplied a (time, frequency) pair as input parameters to the forward harmonic sinusoid tracking algorithm. A third input, the partial index which tells which partial of the event this (time, frequency) pair points to, is supplied in the partial selector on the HS setting tab page. This initial user input becomes the first *anchor atom* of the harmonic sinusoid.

### 8.2.1 Partial selector

Choose a partial index from the list of numbers to indicate the current partial index used as input (along with time and frequency supplied by the mouse pointer) to the harmonic tracking algorithm to select or reselect a harmonic sinusoid.

Tip: One may also click on a sinusoidal atom in the main display to set the partial selector to the atom's partial index.

The selected harmonic sinusoid may have various corrigible or incorrigible errors. HV provides simple interface with which the user can correct the corrigible errors.

### 8.2.2 Algorithmic parameters of harmonic sinusoid tracking

The following six adjustable algorithmic parameters are specified on the HS setting page.

maxp	maximal number of partials, 50 by default
delm	maximal frequency estimate error, in bins, 1 by default
delp	maximal pitch jump, in semitone per frame, 1 by default
maxB	maximal stiffness coefficient, 0.001 by default
minf0, maxf0	fundamental frequency range, 27.5~3520 (A0~A7) by default

These parameters are usually used at their default values. Set delp to 0 if the target is constant-pitched or larger than 1 if there are fast frequency variations. However, in the latter case changing the window size is probably more efficient.

### 8.2.3 Manual termination and continuation

If the current selected harmonic sinusoid is longer than it should be at either end, i.e. it contains extra frames before or after its genuine duration, these extra frames should be

removed from this harmonic sinusoid. Extra frames before the genuine start of the harmonic event are removed by holding down X key and clicking on an atom of the desired starting frame. Extra frames after the genuine end of the harmonic event are removed by holding down X key and clicking on an atom of the desired last frame.

If the current selected harmonic sinusoid is shorter than it should be at either end, i.e. it has not covered the whole duration of the harmonic event, it should be extended to include extra frames before its current start or after its current end. The user can extend the harmonic sinusoid backward from its beginning by clicking an atom of the current starting frame, or extend the harmonic sinusoid forward from its end by clicking an atom of the current last frame.

## 8.2.4 Modifying selected events

The current *selection* of harmonic sinusoid (notice this is not the audio event represented by the harmonic sinusoid) can be modified by dragging existing atoms to desired places in the harmonic selection mode. When an atom is dragged, it is relocated to the spectral *peak* nearest to the pointer position, with immediate update to other atoms at the same frame. When the dragging is finished (i.e. the mouse key released), the whole harmonic sinusoid is updated to the change.

Algorithmically, when an atom is dragged to a spectral peak by the user, the harmonic sinusoid tracker associates a new anchor atom to the harmonic event just the same way as it associates the first anchor atom using the initial input, unless the atom being dragged is already an anchor atom. The new anchor atom must be compatible with previous anchor atoms. Therefore if there are already other anchor atoms at the same frame, the frequency range of the new anchor atom will be limited by them. All anchor atoms are marked by a “+” in a square in the spectrogram display and have the same priority in harmonic sinusoid tracking.

The use of multiple user inputs as anchor atoms helps to provide good harmonic sinusoid selections. However, incompatibility with existing anchor atoms may prevent a new anchor atom being moved to the desired spectral peak. In this case it is necessary to enlarge the algorithmic parameter *deltm*, or remove one or more existing anchor atoms at the same frame. To change an anchor atom back to a normal atom, move the mouse pointer at the atom and press Delete.

Tip: Hold Shift key down during this action to change all anchor atoms at the pointed frame to normal atoms; hold both Shift and Ctrl keys down during this action to changes all anchor atoms to normal atoms.

## 8.3 Editing harmonic sinusoids

Selected harmonic sinusoids represent the deterministic part of harmonic sound events. The operations presented below do not address transients or noise of such events.

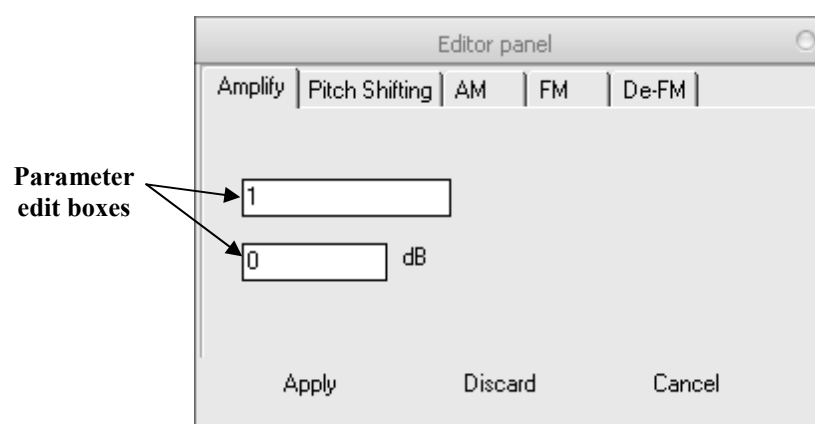
### 8.3.1 Extracting/removing harmonic sinusoids


To extract the selected harmonic sinusoid from audio, move the mouse pointer to any of its atoms (marked in the spectrogram display by “+” tags) and choose Extract from the popup menu. Alternatively, choose Extract from the popup menu of the Event viewer. This operation removes all other concurrent event, as well as the selected event’s own transients, noise components and reverbs. Extract does not alter the harmonic selection itself; repeating the Extract operation yields the same result.

To subtract the selected harmonic sinusoid from the audio, move the pointer to any of its atoms and choose Cut from the popup menu. This operation removes the deterministic part of the selected event but leaves its transients, noise components, and reverbs behind. Cut does not alter the selection itself; repeating the Cut operation once inverts the selected harmonic sinusoid.

### 8.3.2 Editor panel

Other editing functions are available from the Editor panel. To launch the editor panel, move the mouse pointer to an atom of the selected event and choose from the “Editor panel...” submenu. Currently 5 operations are available, each occupying a page of the Editor panel. The operations are triggered typing into or mouse-wheeling the parameter edit boxes. Although the user is free to switch between the pages, it is suggested that only one type of editing operation is carried out during each Editor panel session, marked at either end by the launch or closing of the Editor panel. The combination of different operations during a single Editor panel session has undefined behaviour.



Editing operations via the Editor panel immediately updates the audio content observed in the main waveform/spectrogram display. One can listen to the results using the Play button () or select Play from the popup menu in the main display.

The editor panel is closed using the Apply or Cancel button. If the panel is close using Cancel button, all changes done during the Editor panel session are undone. If the panel is closed using Apply button, both the editing result and the change of sinusoidal parameters are kept. This enabled concatenation of Editor panel sessions for further editing. Use F2 key to restore the original audio content at any time.

### 8.3.3 Amplification

Amplification operation amplifies the selected harmonic sinusoid while leaving the rest of the audio contents intact. Amplification is operated via the Amplify page on the Editor panel. The two editor boxes on this page show the magnitude gain on linear and dB scales, respectively. One can change the gain factor by typing into either editor box and pressing Enter, or mouse-wheeling inside either editor box.

Tip: Mouse-wheeling an atom of the selected harmonic event in the normal selection mode has the same effect as mouse-wheeling inside the linear scale gain edit box. It will also bring up the Editor panel and set it to the Amplify page if it is not already shown.

### 8.3.4 Pitch shifting

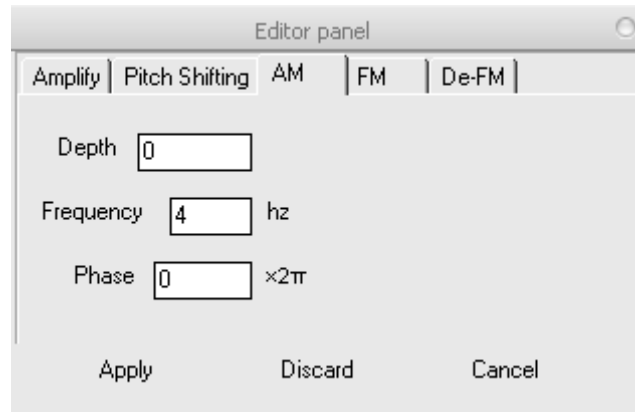
Pitch shifting operation shifts the pitch of the selected harmonic sinusoid while leaving the rest of the audio contents intact. Pitch shifting is operated via the Pitch Shifting page on the Editor panel. The editor box on this page shows the amount of pitch shift in semitones. One can change the pitch by typing into the editor box and pressing Enter, or mouse-wheeling inside the editor box.

Tip: Mouse-dragging an atom of the selected harmonic event in the normal selection mode shifts the pitch to relocate the dragged atom to its new position. It will also bring up the Editor panel and set it to the Pitch Shifting page if it is not already shown.

If the “Duplicate” checkbox on this page is checked, the shifted harmonic sinusoid appears as a duplicate of the original one instead of replacing the latter, creating a chorus effect.

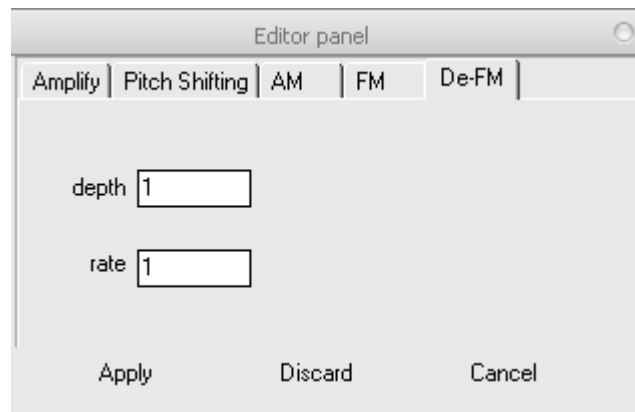
### 8.3.5 Amplitude and frequency modulations

AM/FM operations apply amplitude/frequency modulations to the selected harmonic sinusoid. If the harmonic sinusoid was already modulated, the new modulations are superposed onto the original modulations. AM/FM are operated via the AM/FM pages on the Editor panel. Each of the two pages contains three editor boxes for modulation depth, frequency and phase, respectively. Type into or mouse-wheel inside these editor boxes to change the corresponding modulation parameters.



### 8.3.6 Frequency demodulation and remodulation


Frequency remodulation changes the depth or rate of existing frequency modulations of the selected harmonic sinusoid. Remodulation is operated via the De-FM page on the Editor panel. The page contains two editor boxes for relative modulation depth and rate respectively, both of which are 1 by default. One can change the magnitude or speed of the modulation by typing into the editor boxes and press Enter, or mouse-wheel inside the editor boxes. Frequency demodulation is obtained by setting the relative modulation depth to 0.





### 8.3.7 Saving the editing result

Use “File|Save as” from the main menu or the “Save as” button on the toolbar to save the currently displayed audio content as a separate file. To save the whole audio, first zoom-out use “Restore|Time zoom” to display the whole duration, then save.



Changes to the harmonic sinusoidal model parameters are saved using the Save button on the Event viewer.

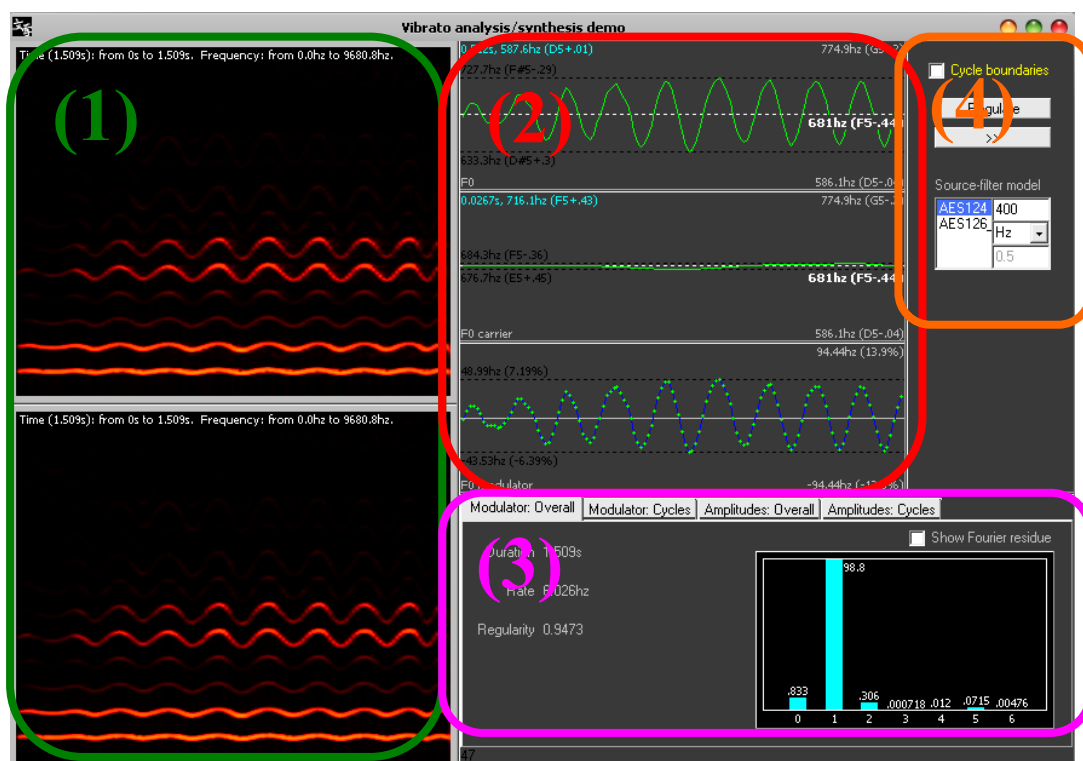
## 8.4 Vibrato analysis and synthesis demo

HV implements the vibrato representation of harmonic sinusoids and demonstrates the use of this representation for the analysis and synthesis of vibratos. To launch the demo, choose Vibrato... from the popup menu of the Event viewer, or alternatively hold down the Shift key and double-click the event to analyze in the Event viewer.

Tip: The vibrato analysis/synthesis demo only makes sense if the selected event contains regular vibratos. To analyze irregular vibratos with locally regular sections, it may be sensible to analyze a regular section only by removing its leading and trailing frames (see 8.2.3 for how this is done).

Top: Re-launch the vibrato analysis/synthesis demo to discard any changes made to the event via the demo interface.

### 8.4.1 GUI



The GUI of the Vibrato demo consists of two spectrogram display panes (1), three pitch trajectory display panes (2), the vibrato representation display/editing panel (3), and control panel (4) with a few other controls.

When a harmonic sinusoid is dispatched to the Vibrato demo by launching it, HV makes a copy of the original audio content over the duration of the harmonic sinusoid and send it to the upper spectrogram pane. This copy contains the harmonic sinusoid as well as other concurrent events. The harmonic sinusoid itself is displayed in the lower spectrogram pane.


During the Vibrato demo session the audio content displayed in the lower pane can be changed via the demo GUI. Basic zoom-in, zoom-out and playback functions are available from these spectrogram panes via mouse-wheeling and the popup menu.

The three pitch trajectory panes, from top to bottom, respectively show the trajectories of the whole pitch, its carrier part and its modulator part. The lower panes can receive user input for pitch shifting and remodulation.

The vibrato representation panel provides detailed information on the parameterization of the harmonic sinusoid using the vibrato representation over four tab pages. The “Modulator:Overall” and “Amplitudes:Overall” tab pages provide frequency and source-filter parameterizations averaged over the whole duration; one can change these parameters to obtain vibrato modification on global scale. “Modulator Cycle” and “Amplitudes Cycle” tab pages provide cycle-by-cycle frequency and source-filter parameterizations; one can change these parameters to modify vibrato parameters at individual cycles.

For details on the vibrato representation of harmonic sinusoids, see [1].

## 8.4.2 The pitch trajectory panes




The Vibrato demo contains three pitch trajectory panes. The F0 pane (top pane) shows the current pitch trajectory of the harmonic sinusoid shown in the lower spectrogram pane. The vibrato representation parameterizes a harmonic sinusoid by modulation cycles. If a distinctive periodical pitch modulator exists, the periodical pitch peaks are taken as the cycle boundaries. They are made visible in the F0 trajectory pane as yellow vertical lines by checking the “Cycle boundaries” checkbox on the control panel. The F0 pane does not accept inputs.

The F0 carrier pane (middle pane) shows the slow-varying part (carrier) of the pitch trajectory in the F0 pane. Although the F0 carrier is not parameterized in the vibrato representation, HV may associate a few control points to it, which are visible as red crosses if the “Cycle boundaries” checkbox is checked. Mouse-wheeling the F0 carrier pane pitch-shifts the F0 carrier (therefore the global pitch). Mouse-dragging the control points of the F0 carrier changes the shape of the F0 carrier, allowing additional long-term pitch modulation.

The F0 modulator pane (bottom pane) shows the periodically varying part (modulator) of the pitch trajectory in the F0 pane. Move the mouse pointer over this pane highlights individual cycles and displays the vibrato representation of the highlighted cycle in the vibrato representation panel. Move the pointer before the first or after the last complete cycle to cancel the highlighting and display the overall (average) vibrato representation. Mouse-wheeling the F0 modulator pane changes the modulator extent either of a single cycle (if one is highlighted), or of the whole vibrato (if none is highlighted).

Tip: Hold Ctrl/Shift down for coarser(faster)/finer(slower) parameter adjustment via mouse-wheeling in the Vibrato demo.

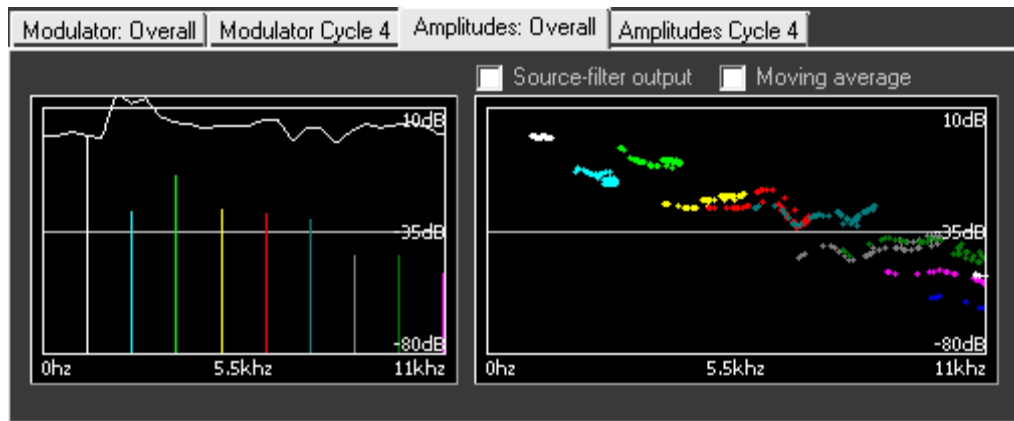
### 8.4.3 The vibrato representation panel



The vibrato representation panel displays parameters of the vibrato representation of the harmonic sinusoid, and allows modification to the harmonic sinusoid by changing these parameters. It contains four tab pages: two labelled Modulator and two labelled Amplitude. The Modulator pages handle the parameterization of the frequency, and the Amplitude pages handle that of the amplitude. In each of the two pairs one tab is labelled Cycle# for showing and changing single-cycle parameters and the other tab is labelled Overall for showing and changing the average parameters across all cycles. Moving the mouse pointer in the F0 modulator pane tells HV which cycle two Cycle# pages currently handle.

The Modulator:Overall page shows the duration of the harmonic sinusoid, the modulation rate and the modulator regularity, and in a histogram the Fourier series decomposition of the average modulator cycle (magnitude only, without phase) using the modulation rate as fundamental frequency. The duration, modulation rate can be modified by mouse-wheeling the values displayed; the modulator shape can be modified through the Fourier series coefficients by mouse-wheeling the histogram.

The Modulator:Cycle# page shows the duration, modulation rate, average pitch, and modulation extent/depth of the concern modulation cycle, as well as the Fourier series histogram of the modulator cycle using the cycle modulation rate as the fundamental frequency. Notice that in this case the duration and modulation rates are a pair of reciprocals, and one can change the cycle duration by mouse-wheeling either of them. The modulator shape over this cycle can be modified through the Fourier series coefficient by mouse-wheeling the histogram.




The Amplitudes:Overall page contains two charts. The left chart shows the source-filter model of the harmonic sinusoid, in which the source model (partial amplitudes) are plotted as vertical bars at the partial frequencies, each partial in a different colour, and the filter model is plotted as a curve of dB gains against frequency. The right chart shows the observations of partial amplitudes, in which each observation (i.e. amplitude at one frame) is plotted as one “+”, coloured in accordance with the source-filter chart, in the (frequency, dB-amplitude) plane. If the “Source-filter output” checkbox is checked, the plotted amplitudes are observed through the source-filter model; if not, the plotted amplitudes are raw observations. In the latter case one can tick the “Moving average” checkbox to suppress observation noise. One can change the source model by dragging or wheeling either chart, but cannot change the filter model.

The Amplitude:Cycle# page has the same arrangement as the Amplitude:Overall page but is dedicated to the source-filter modelling over single cycles. The Vibrato demo imposes identical filter curve over all cycles but allow the source model to vary from cycle to cycle. These single-cycle source models can be changed by dragging or wheeling the charts.

## 8.4.4 Source-filter modelling settings

The source-filter modelling techniques used in the Vibrato demo include the original version native to the vibrato representation (AES124 [1]) and a later version (AES126 [2], slow-variation method) adapted to work with the vibrato representation. One can specify which technique to use by choosing from the list (AES124/AES126) on the control panel. Other available settings include the width of frequency bands used by the filter model on Hz or mel scale, as well as a balancing factor used by the slow-variation method in [2].



## 8.4.5 Other controls

Use the Regulate button to apply the average modulation cycle to the whole harmonic sinusoid. This demonstrates the removal of modulator noise due to inaccurate measurement or bad playing/singing technique, although a more delicate approach should be taken in practical cases.

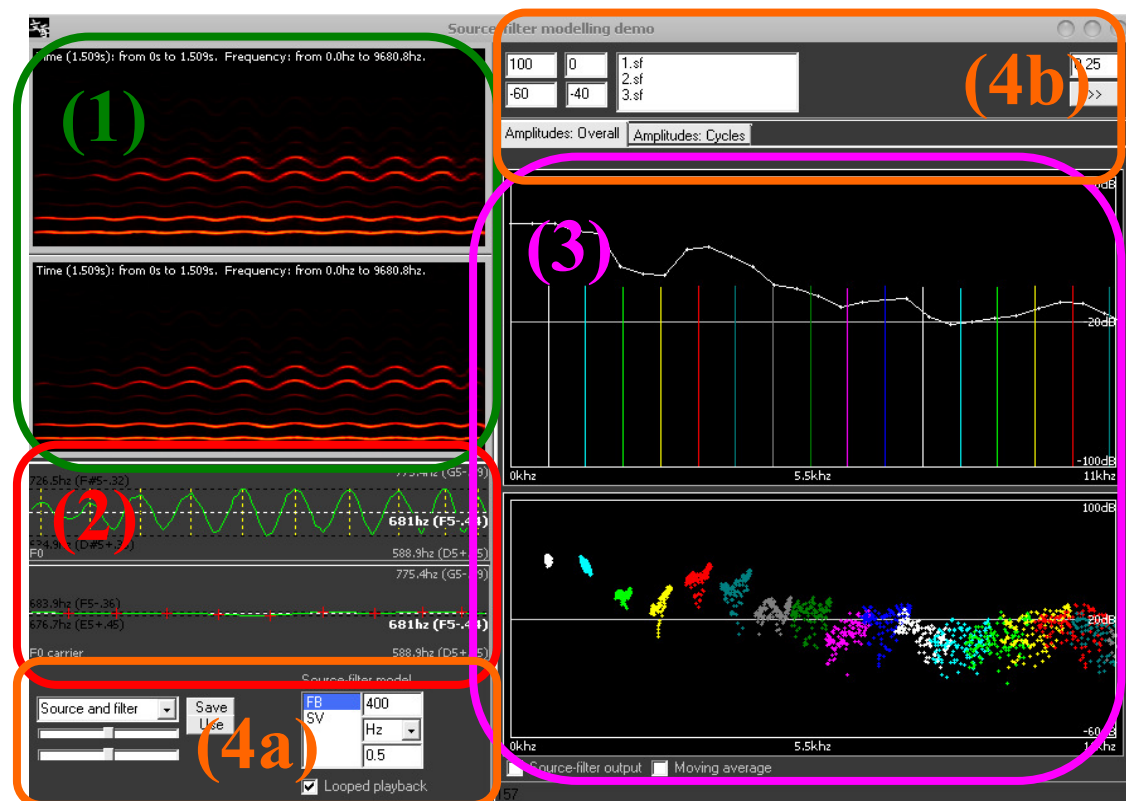
Use the “>>” button to send the modified harmonic sinusoid back to the main display, replacing the original harmonic sinusoid with which the Vibrato demo was launched. If the duration of the harmonic sinusoid is stretched, only the beginning within the original duration is sent.

## 8.5 Source-filter modelling demo

HV implements the source-filter representation of harmonic sinusoids and demonstrates the use of this representation for timbre handling of harmonic sinusoids. Although the vibrato representation also handles source-filter modelling, the source-filter representation improves on that by fine-grained frame-by-frame modelling, parameterizing the filter model and allowing it to vary in time, and provide the flexibility of arbitrary segmentation (i.e. not limited to modulation cycle boundaries). To launch the demo, choose Source-fitter... from the popup menu of the Event viewer, or alternatively double-click the event to analyze in the Event viewer.

Top: Re-launch the Source-filter modelling demo to discard any changes made to the event via the demo interface.

### 8.5.1 GUI



The GUI of the Source-filter demo consists of two spectrogram display panes (1), two pitch trajectory display panes (2), the source-filter representation panel (3), and two control panels (4a and 4b).

The two spectrogram panes are arranged the same way as in the Vibrato demo. To play sounds use the popup menus of these spectrogram panes. The two pitch trajectory panes handles the F0 modulator (upper pane) and F0 carrier (lower pane) respectively. The source-filter representation panel contains two tab pages for single-segment and overall-average source-filter models respectively. One can move the mouse pointer inside the F0 modulator

pane to specify which segment is currently handled by the single-segment source-filter model tab page. Each tab page displays the source-filter model in the upper chart and the amplitude observations in the lower chart, the same way they are displayed in the Vibrato demo. The control panels hosts several algorithmic and display settings and a few other controls.

For details on the source-filter representation of harmonic sinusoids, see [2].

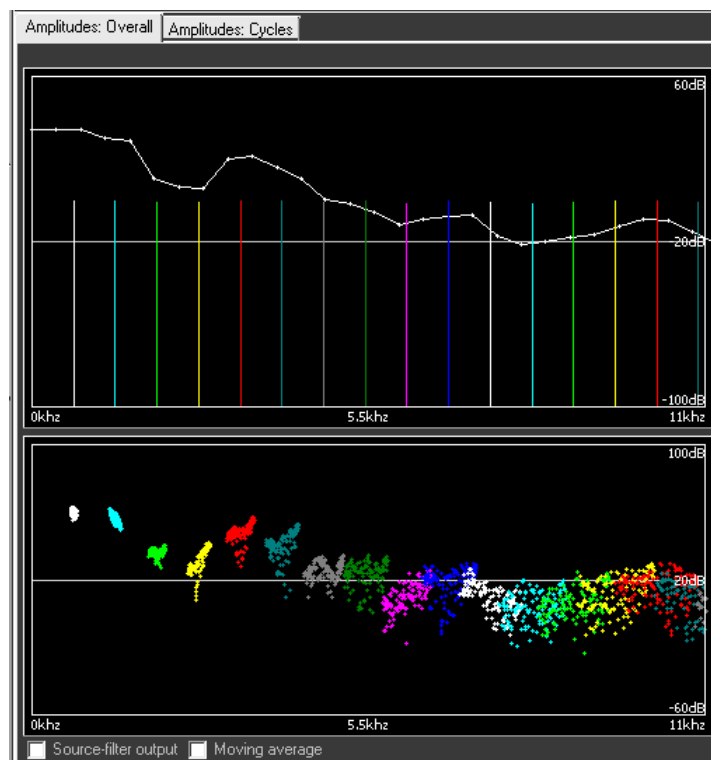
### **8.5.2 The pitch trajectory panes**

The Source-filter demo handles pitches by segments, which are similar to modulation cycles used in the Vibrato demo. When a harmonic sinusoid is dispatched to the Source-filter demo by launching it, HV tries to find a periodical pitch modulation to slice the whole duration into segments, and displays to segment boundaries as vertical bars in the F0 modulation pane. If the harmonic sinusoid has distinctive vibrato, then the segments are the same as the modulation cycles. Unlike the vibrato demo, this segmentation does not take part in the source-filter analysis, but is only used to allow source model modification on a group of adjacent frames, and to provide F0 carrier control points. To insert a new segment boundary, hold Shift key down and click at the point a new boundary is to be added. To remove an existing segment boundary, hold Shift key down and click on the boundary. To move a segment boundary to another position, use the mouse pointer to drag it.

For each of the segments a F0 carrier control point is assigned and plotted in the F0 carrier pane as red crosses. These control point can be dragged with the mouse pointer to modify the F0 carrier. One can also pitch-shift the whole F0 carrier (therefore the harmonic sinusoid) by mouse-wheeling the F0 carrier pane.



### 8.5.3 The source-filter representation panel



The source-filter representation The Amplitudes:Overall tab page on the source-filter representation panel plots the average source-filter model in the upper chart and complete amplitude observations in the lower chart; the Amplitude:Cycles tab page plots single-frame source-filter models in the upper chart and single-segment amplitude observations in the lower chart, in which the frame or segment is specified by moving the mouse pointer inside the F0 modulator pane. In the source-filter model charts the source model is plotted in coloured vertical bars at the partial frequencies, and the filter model is plotted as a curve of dB gains against frequency. If the “Source-filter output” checkbox is checked, the plotted amplitudes are observed through the source-filter model; if not, the plotted amplitudes are raw observations. In the latter case one can tick the “Moving average” checkbox to suppress observation noise.

Modifications to the filter model are available only on the Amplitude:Overall tab page on a global scale via the average filter model. The source-filter model chart offers filter control points, whose frequencies are specified by the source-filter model settings. These control points can be dragged to other gain levels to change the average filter model, while the same amount of change is applied to all the actual frame-wise filter models. If the Shift key is held down during this operation, all control points with higher frequency than the one being dragged are moved along. If the Ctrl key is held down during this operation, the average source model (and the frame-wise source models) is updated in the counter direction by the same amount.


Modifications to the source model are available on both tab pages, although manually modifying frame-wise source model is hardly useful. The average source model is modified by dragging or mouse-wheeling inside either the average source-filter model chart or the

complete amplitude observations chart. The frame-wise source model is modified by dragging inside the frame-wise source-filter model chart or the single-segment amplitude observations chart. Notice that due to smooth synthesis the source model modification of one frame may also affect the source models of nearby frames.


Tip: The following additional functionalities are available during the average source model modification by dragging: a) if Shift key is held down then all harmonic partials of the dragged one are also modified by the same amount; b) if a number key (1~9) is held down then all higher partials whose partial index differs from the one being dragged by a multiple of that number are also modified by the same amount.

## 8.5.4 Source-filter modelling settings

The source-filter modelling techniques used in the Source-filter demo include the filter-bank (FB) and slow-variations (SV) methods [2]. One can specify which technique to use by choosing from the list on control panel (4a). Other available settings include the width of frequency bands used by the filter model on Hz or mel scale, as well as a balancing factor used by the slow-variation method.

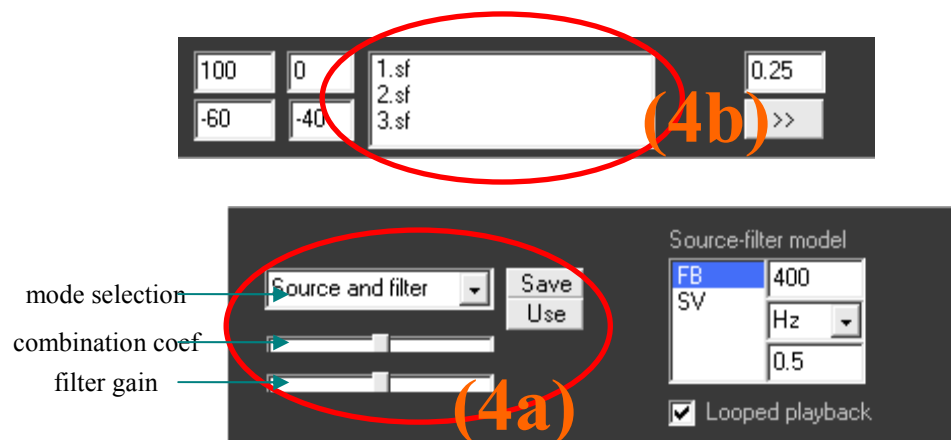


## 8.5.5 Display settings



Five edit boxes control the dB ranges and frequency ranges to plot the source-filter models and amplitude observations. The two amplitude dB-range boxes specify the extents of the y axis in the amplitude observations charts. This range is shifted by the amount specified by the source-filter dB-range shift box to specify the extents of the y axis in the source-filter model charts. If necessary, the filter curve can be shifted by the amount specified in the filter shift box. The frequency range displayed in the charts is specified by the cut-off frequency box on the discrete frequency scale (sampling frequency=1). All these five are merely display settings and have no effect on source-filter analysis or synthesis.

## 8.5.6 Saving and using source-filter models



A source-filter model can be saved using the Save button on control panel (4a). By default source-filter models are saved in the same directory as the programme executable with file extension .sf. The saved models appear in a list on control panel (4b) labelled by their file names. To save a source-filter model with a user-chosen file name, hold Shift key down when pressing the Save button.

Previously saved source-filter models can be applied to the harmonic sinusoid the Source-filter demo is currently working with. To do this HV constructs a destination source-filter model as a weighted average of the current and the stored source-filter models, in which the combination coefficient is controlled by the combination coef. floater bar on control panel (4a). Setting the floater to the leftmost indicates no change (destination source-filter model is identical to current model); setting the floater at the middle indicates full change (destination source-filter model is identical to stored model); setting the floater to the rightmost indicates double full change. The destination filter model can be shifted up to  $\pm 10$ dB using the filter gain floater bar.

The mode selection box provide five ways to apply the destination model: applying both source and filter models, applying source model only, applying filter model only, smooth slide towards the destination model, and smooth slide to-and-fro the destination model. The actual application of the destination model is invoked by pressing the Use button or by dragging the floater bars below the mode selection box. The change is applied to the source-filter model as well as the audio. To undo the change re-launch the Source-filter demo.

## 8.5.7 Other controls

Use the “Looped playback” checkbox to enable/disable playing the audio content in the spectrogram display pane in loops.

Use the “>>” button to send the modified harmonic sinusoid back to the main display, replacing the original harmonic sinusoid with which the Source-filter demo was launched.

## Further reading

Wen X. and M. Sandler, “Sinusoid modelling in a harmonic context,” in *Proc. DADx’07*, Bordeaux, 2007.

## References

[1] Wen X. and M. Sandler, “Analysis and synthesis of audio vibrato using harmonic sinusoids,” in *Proc. AES 124th Convention*, Amsterdam, 2008.

[2] Wen X. and M. Sandler, “Source-filter modelling in sinusoid domain,” in *Proc. AES 126th Convention*, Munich, 2009.

END OF FILE■