

HARPEX-X

version 1.6

manual

Licensing

It is necessary to register with a name and email address to download the plug-in. Each plug-in is associated with a license number. The license number is part of the download link, which will be sent to you by email. This license number is also imprinted in the plug-in along with the name of the licensee at the time it is downloaded.

Trial period

During the trial period, the plug-in contacts a license server each time it is loaded, in order to verify that the trial period has not expired. A network connection is therefore necessary during the trial period. If you choose to purchase the plug-in, you may activate up to two computers at a time.

Activation

The activation process itself requires a network connection. Once a computer is permanently activated, it will not contact the license server anymore, and therefore does not require a network connection. If you upgrade or reinstall the plug-in, it will need to be reactivated the next time you use it. Reactivation does not affect your activation counter. Different flavors of the plug-in (32 bit, 64 bit, AAX, VST, Audio Unit) may also require separate activation, which will not affect your activation counter as long as they are installed on the same computer.

Transfer

After activating two computers, you may transfer one of the activations to a third computer. To do this, install the plug-in on the new computer. The first time you use it, you will be given the option of deactivating one of the previously activated computers. The first transfer may take place any time. Subsequent transfers are limited to one per year.

Temporary activation

During activation, you will be asked if you want to activate your computer permanently or temporarily. If you need to move the plugin between different computers more often than once per year, you have the option of activating your computers temporarily. This allows you move the plugin once per day, but requires an internet connection for use.

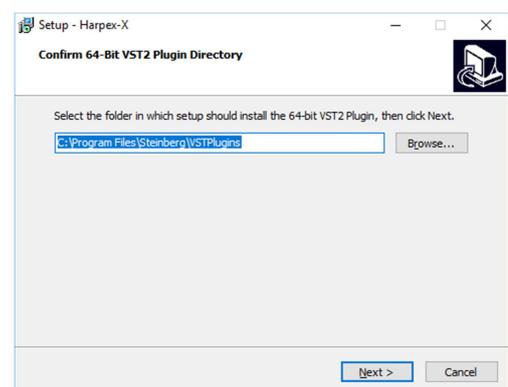
Upgrading / re-installation

To upgrade or download additional copies of the plug-in, use the same download link as the first time, i.e. do not register again, since your activations are linked to the license number, which appears in the download link. Whenever you download, you get a fresh copy of the latest available version.

Installation

Windows

An installation program, setup.exe, is provided. This program will install this manual and copy the plugin to the known



locations for those VST/AAX hosts that are recognized and installed on your computer. All copies created during installation will be removed when you uninstall the plugin via the normal software uninstallation procedure. A list of recognized VST/AAX hosts is displayed during the installation process.

If your VST/AAX host is not on that list, and it does not use the common VST directory (commonly C:\program files\VSTplugins), you must either copy the file Harpex-X.dll from the installation directory (commonly C:\program files\Harpex-X) to the hosts's plugin directory; or add the installation directory to your host's plugin search path.

Mac

The plugin is distributed in a disk image containing a VST, Audio Unit and AAX version of the plugin. To install the plugin, simply drag the desired version (or versions) into the folder aliases provided in the disk image.



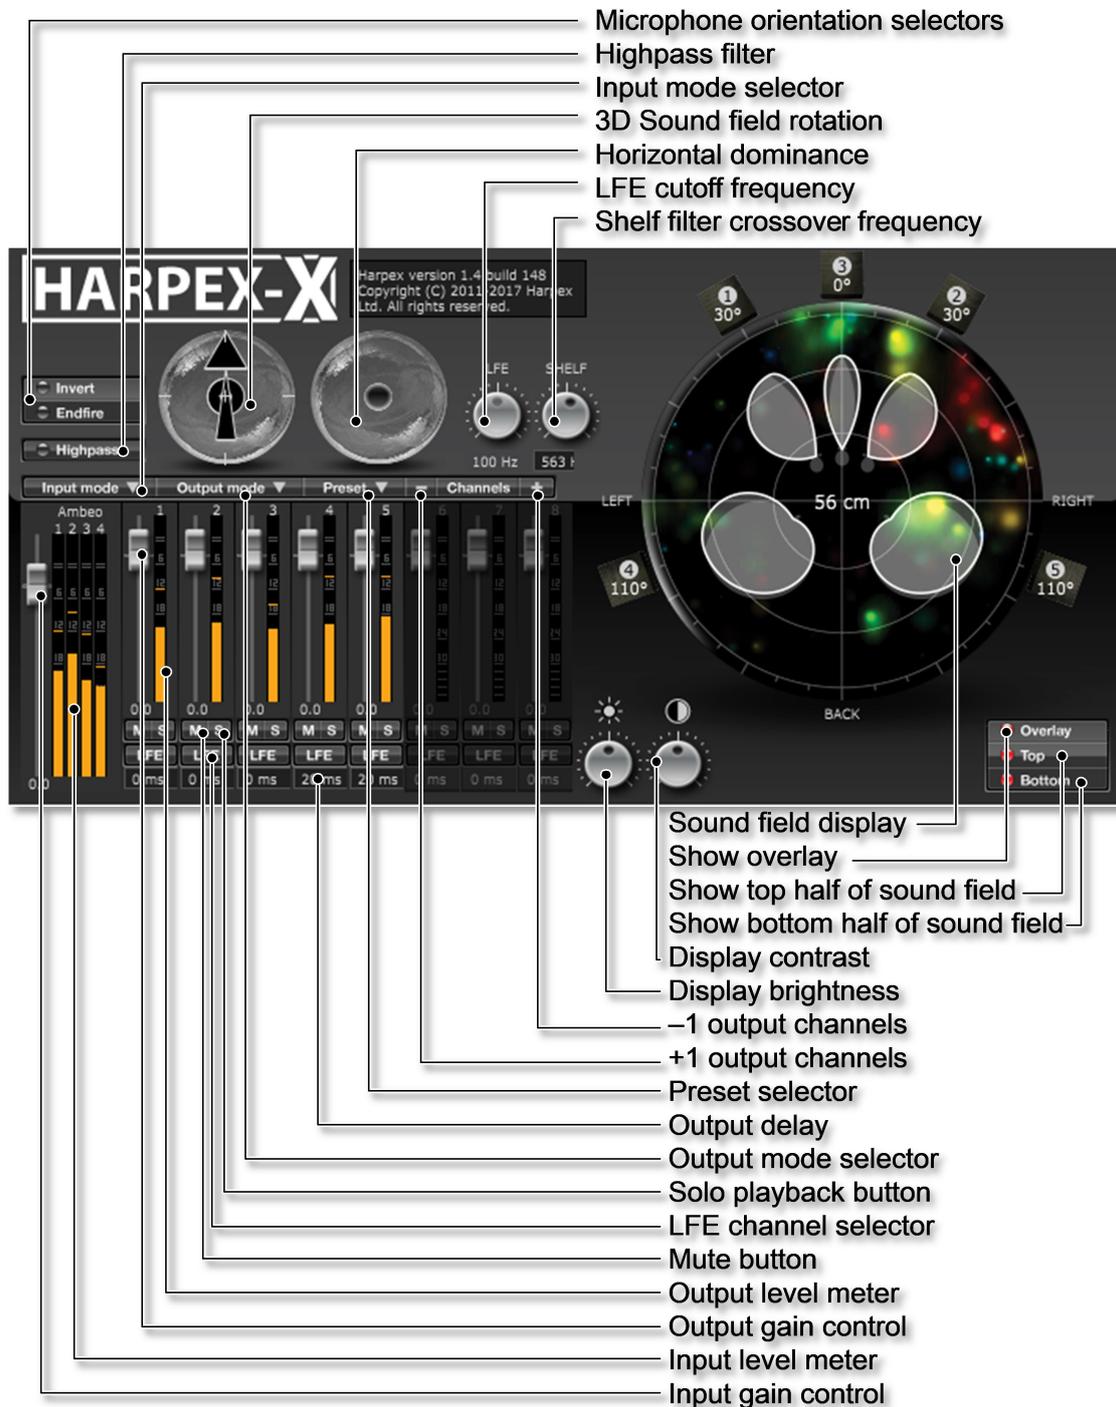
DAW integration

Once the Harpex-X plugin is installed, it should be available as an effect in your Digital Audio Workstation. Since the input to Harpex is a multi-channel signal, it only makes sense to apply it to a multi-channel track. The track may contain any number of channels, but only the first four are used.

For surround decoding, it is important to check the output channel routing. The channel numbers shown in the user interface correspond to the order in which the channels are sent to the DAW. However, some DAWs rearrange these channels before the output is sent to an output device or bounced down to a sound file, so it is essential to check that the channel routing before use.

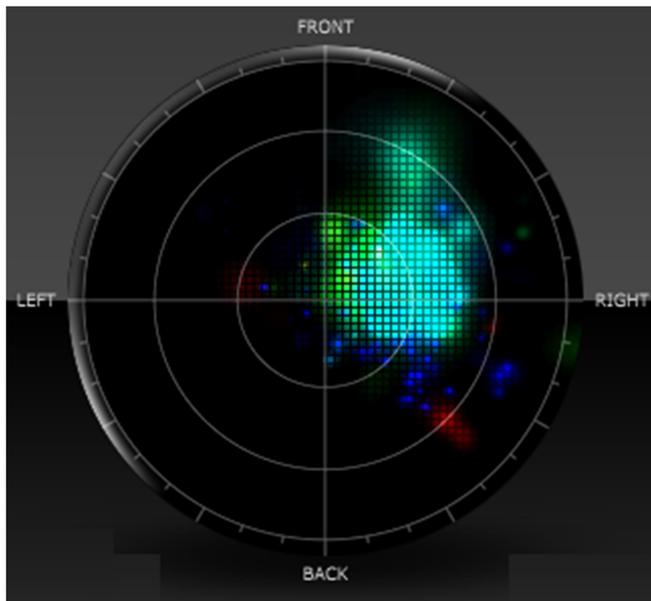
Operation

The Harpex-X plugin contains a number of usage modes that are suitable depending on the desired output audio format. Common to most usage modes are the following controls:



- **Microphone orientation selectors:** Selects which way the microphone was oriented during recording. Used with A-format input.
- **Highpass:** Switch on a highpass filter to reduce wind and handling noise.
- **Input mode selector:** Selects an input mode; FuMa, AmbiX, Double MS or one of the supported tetrahedral microphone arrays.

- **3D Sound field rotation:** Rotates the input ambisonic signal before any other processing takes place. Double-click to reset. See section “Sound field rotation”.
- **Horizontal dominance:** Pulls the entire sound field towards a chosen direction. Double-click to reset.
- **LFE cutoff frequency:** Sets the crossover frequency for the LFE channel and the other channels.
- **Shelf filter crossover frequency:** Sets the crossover frequency between low-frequency and high-frequency panning rules.
- **Sound field display:** Displays the sound field. See section “Sound field display”.
- **Show overlay:** Shows/hides the mode-specific information overlaid on top of the sound field display.
- **Show top half of sound field:** Shows/hides the top half of the sound field in the sound field display. See section “Sound field display”.
- **Show bottom half of sound field:** Shows/hides the bottom half of the sound field in the sound field display. See section “Sound field display”.
- **Display contrast:** Changes the displayable sound field intensity range.
- **Display brightness:** Changes the display brightness corresponding to any given sound field intensity.
- **-1 output channels:** Reduces the number of output channels (only available in some output modes)
- **+1 output channels:** Increases the number of output channels (only available in some output modes)
- **Preset selector:** Selects one of several mode-specific presets (only available in some output modes)
- **Output delay:** Selects the delay per channel in the range 0-100 ms. Used in presets for surround channels.
- **Output mode selector:** Selects an output mode. See section “Output modes”.
- **Solo playback button:** Isolates one output channel or a group of output channels and mutes the others.
- **Mute button:** Mutes the output of one or several output channels.
- **Output level meter:** Displays the instantaneous and peak level per output channel. Hover mouse pointer to read out peak level. Click on level meter to reset peak indicator.
- **Output gain control:** Control output level. In some output modes, output levels can be controlled per channel. Double-click to reset.
- **Input level meter:** Displays the instantaneous and peak level per input channel. Hover mouse pointer to read out peak level. Click on level meter to reset peak indicator.
- **Input gain control:** Control input level. Double-click to reset.



Sound field display

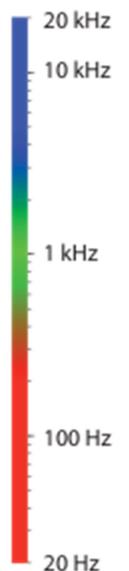
The sound field display shows the spatial and spectral distribution of the signal after rotation and zooming. Different sources of sound will show up as illuminated points or regions in the diagram. Point-like sources are shown as points and sources with an extent are shown as diffuse regions of light. The intensity of the sound corresponds to the intensity of the light. However, since the dynamic range of our hearing is much

greater than what can be represented in the intensity of a pixel, dynamic range is also represented by the size of the region. The size of a spot can therefore correspond either to its intensity or its physical extent. To distinguish, one may observe that extended sound objects are usually of lower visual intensity than point-like sources.

The color of the spots correspond to their frequency content, red is for low frequencies, green for mid-range and blue for high frequencies.

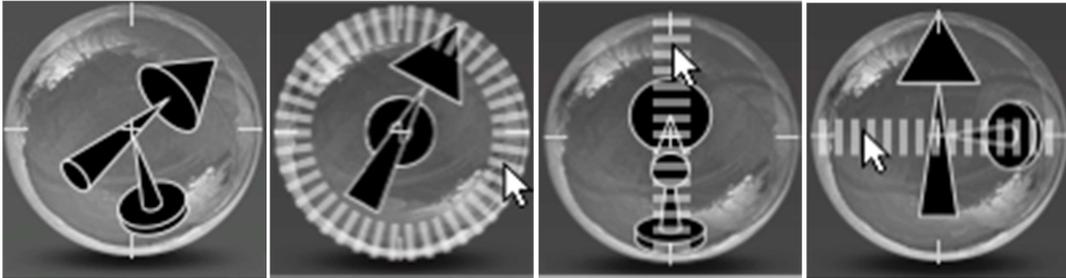
The top of the diagram represents sound coming from straight ahead, the bottom shows sound coming from behind and left and right correspond to left and right. Sound coming from straight above is displayed in the center of the diagram. Sound coming from straight below is also displayed in the center of the diagram. To distinguish between these directions, sound from below is painted with a fine square grid, whereas sound from above is painted in solid color. The illustration above shows a sound from below. It is also possible to use the “Show top/bottom half of sound field” to switch off either the top or bottom half of the sound field.

The Lambert azimuthal equal-area projection is used to project the sphere onto the circular plane. The concentric circles in the diagram represent 0° , $\pm 30^\circ$ and $\pm 60^\circ$ elevation.



Sound field rotation

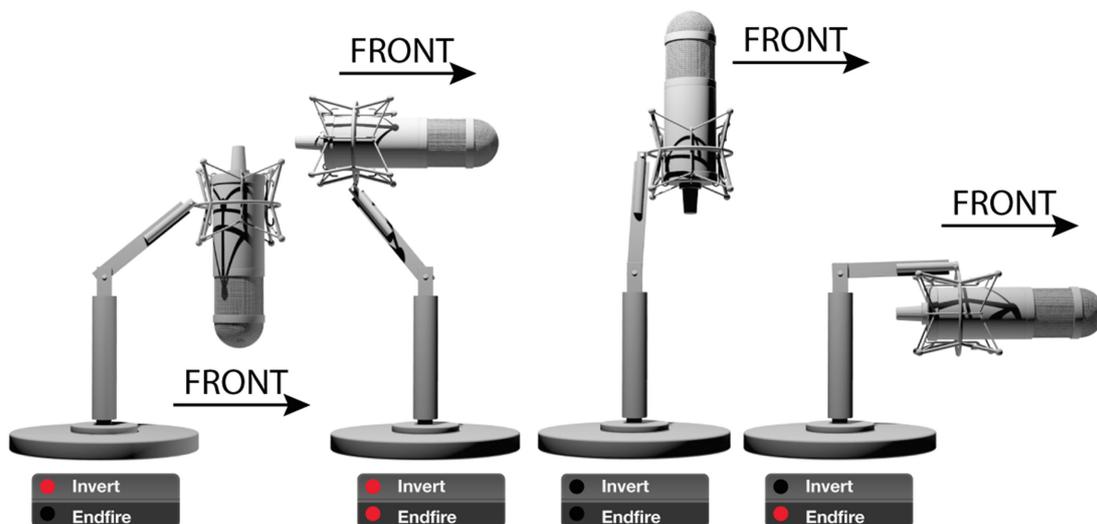
An ambisonic signal can be rotated in any direction. The rotation is represented by an arrow inside a bubble. The arrow represents the attitude of the observer. It has a foot to distinguish between the up and down directions. When this arrow points directly towards the top of the screen and its foot is shown directly behind it, the field is not rotated. To rotate the field, click on a point on the bubble and drag that point to the desired location.



For extra convenience, rulers will appear when the mouse pointer hovers over one of the three areas as shown above. Use the circular ruler around the periphery to rotate the sound field in the horizontal plane. This is the most common use of the rotation bubble. Note that to shift the sound field to the left, you must rotate the arrow to the right, as it represents the orientation of the virtual microphone, not of the sound field. To tilt the sound field up or down, use the vertical ruler. This is useful to shift between endfire and side-address configurations and to compensate for a height difference between the sound sources and the microphone in case the physical microphone was mounted vertically. To complete the set of rulers, there is a horizontal ruler which can be used to rotate the virtual microphone sideways.

Preset orientations

When using a raw tetrahedral input, it may be useful to select one of these preset microphone orientations:



Input modes

- **FuMa:** W, X, Y and Z. The Z channel is optional
- **AmbiX:** The channel order is 1, 2, 3, 4.
- **Double MS:** The channel order is Front, Side, Back, and Z. The Z channel is optional, and not usually present in a double MS signal. The Side and Z channels are amplified by 2.7 dB compared to the Front and Back channels. I.e. a dry source moving around the array should reach a peak level 2.7 dB higher in the Front and Back channels than in the Side channel.
- **Raw tetrahedral:** The inputs named “Ambeo”, “H3-VR”, “NT-SF1”, “SPS-200” and “TetraMic” are the raw microphone signals from tetrahedral microphone arrays. The channel order for all of these microphones is
 - 1: Front left up
 - 2: Front right down
 - 3: Back left down
 - 4: Back right up

Output modes

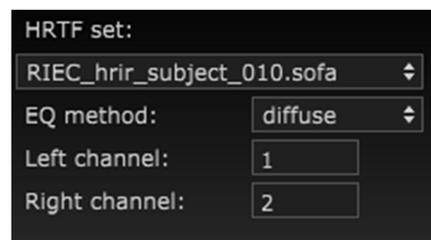
The following output modes are available:

- **Binaural:** Use this for headphone playback.
- **Stereo:** Use this for playback over stereo loudspeakers.
- **Shotgun:** Use this to isolate sound sources based on their spatial location or to play back over custom 3D loudspeaker arrays.
- **Surround:** Use this for playback over a horizontal loudspeaker rig.
- **AmbiX:** Enables the use of the plugin as an ambisonic processor and upsampling to higher-order ambisonics.
- **Speakerdome:** Decode to one of several standard 3D loudspeaker array layouts.

Binaural

This output mode is most suitable for headphone playback. There are two mode-specific settings:

- **HRTF set:** Harpex-X contains a database of several measured HRTF sets. The default setting is KU100, which is measured on a dummy head designed to provide the best possible reproduction for the largest possible audience. Select “Other...” to load an HRTF set from a file. The file should be in the SOFA format.
- **EQ method:** HRTFs add different coloration to sounds coming from different directions. This control lets you select from which direction sound should pass without coloration:
 - **Diffuse:** Sound emanating equally much from all directions.
 - **Horizontal:** Sound emanating equally much from all horizontal directions.



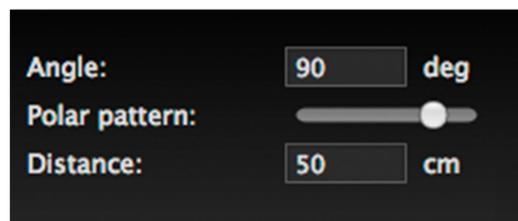
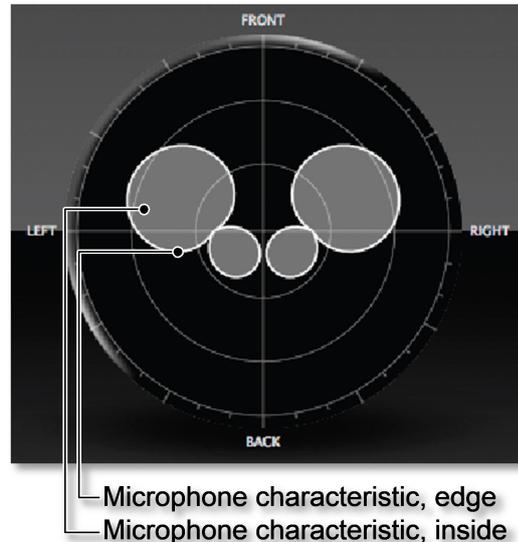
- **Front:** Sound directly from the front
- **None:** Uses the equalization inherent in the measured impulse responses.
- **Left channel / right channel:** In some circumstances, the left and right channels are not assigned numbers 1 and 2. To simplify the routing in the DAW, you can reassign the channels here.

The HRTF sets that are built into Harpex are provided by the IRCAM Listen project and the SADIE project at the University of York, respectively.

Stereo

This output mode is most suitable for playback over stereo loudspeakers. There are three mode-specific settings:

- **Angle:** Selects the angle between the pair of virtual microphones. The angle can be set in the angle input box or by dragging the edge of the microphone characteristic.
- **Polar pattern:** Selects the polar pattern of the microphones, on a scale from omnidirectional via subcardioid, cardioid, supercardioid and hypercardioid to figure-of-eight. The polar pattern can be set with the polar pattern slider or with the mouse wheel when the mouse is over the microphone characteristic.
- **Distance:** Selects the distance between the virtual microphones. Changing the distance does not alter the relative amplitudes of sound sources, but will add direction-dependent phase shifts similar to those in real-world spaced microphone techniques. The distance can be set in the distance input box or by dragging the inside of the microphone characteristic.



The following presets are available in this mode:

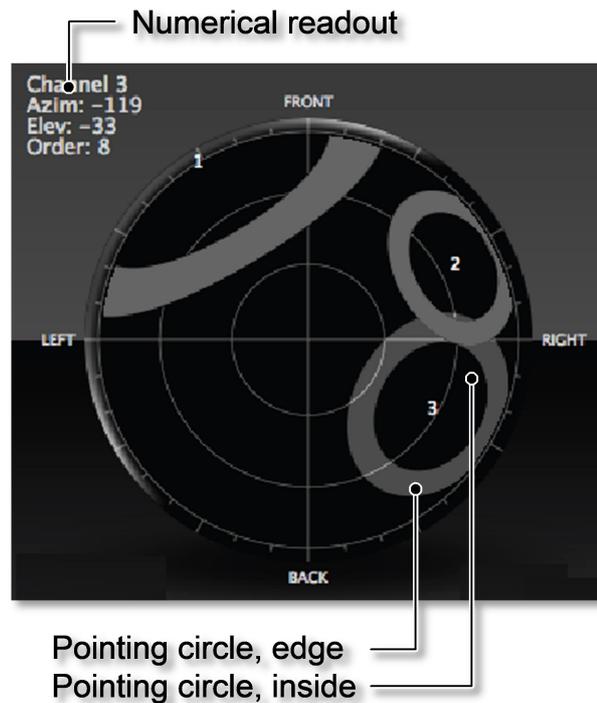
- **XY:** Coincident cardioid microphones at 90° angle.
- **Blumlein:** Coincident figure-of-eight microphones at 90° angle.
- **AB:** Spaced omnidirectional microphones.
- **ORTF:** Spaced (17 cm) cardioid microphones at 110° angle
- **DIN:** Spaced (20 cm) cardioid microphones at 90° angle
- **NOS:** Spaced (30 cm) cardioid microphones at 90° angle

Shotgun / custom 3D loudspeaker array

This output mode is designed to isolate sound sources based on their spatial location. Up to eight output channels can be activated with the +/- channel buttons.

Each channel is associated with a virtual microphone. The direction and directivity of each virtual microphone is indicated with a pointing circle overlaid on the sound field display.

The direction of each microphone can be set by clicking inside its pointing circle and dragging it to its new location. Pointing circles may overlap. If you click on a region covered by more than one circle, you will activate the one whose center is closest to the mouse pointer. In some cases, several virtual microphones are displayed directly on top of each other. This may happen if they point in the same direction or in directions that are symmetrical with respect to the horizontal plane. In such cases, it may be necessary to temporarily mute one of the channels. This will hide the corresponding pointing circle, giving access to the other one. The numerical readout indicates which channel will be affected if you click and drag from the mouse pointer's current position.



Note that microphone directions are, like all other parameters, possible to automate, thereby making it possible to track moving sources.

The directivity of each microphone can be set by clicking on the edge of its pointing circle and dragging it, or by turning the mouse wheel while the mouse pointer is inside the pointing circle. Note that this action may cause the pointing circle to constrict so that the mouse pointer ends up outside the circle, causing the mouse wheel to become inactive. The inner edge of a pointing circle represents the angle where the virtual microphone's gain is -3 dB. The outer edge represents -6 dB. A smaller circle corresponds to a more directive virtual microphone.

It is important to distinguish between upward- and downward-pointing virtual microphones. In the illustration, microphone #2 points upward and #3 points downward. Microphone #1 points in a horizontal direction. An upward-pointing microphone can be moved to the bottom half by dragging it beyond the edge of the sound field display. A downward-pointing microphone can be moved to the upper half in the same manner. See section "sound field display" for a visual method of determining whether a sound comes from above or below.

The shotgun mode can also be used to decode to custom 3D loudspeaker layouts. To decode to a standard 3D loudspeaker layout (Atmos, Auro 3D etc.), see the Speakerdome output mode. In this usage, each pointing circle should point in the direction of the corresponding loudspeaker. The directivity of the virtual microphones should be adjusted so that the pointing circles just overlap. The preset loudspeaker layouts can be used for reference. The following presets are available:

- **Octahedron:** Six loudspeakers arranged like the vertices of an octahedron.

1	2	3	4	5	6
Front left	Front right	Rear left	Rear right	Above	Below

- **3D7.0:** Seven loudspeakers forming a modified 7.0 layout.

1	2	3	4	5	6	7
High front left	High front right	Center front	High rear	Low front	Low rear left	Low rear right

- **Cube:** Eight loudspeakers arranged like the vertices of a cube.

1	2	3	4	5	6	7	8
Low front left	Low rear left	Low rear right	Low front right	High front left	High rear left	High rear right	High front right

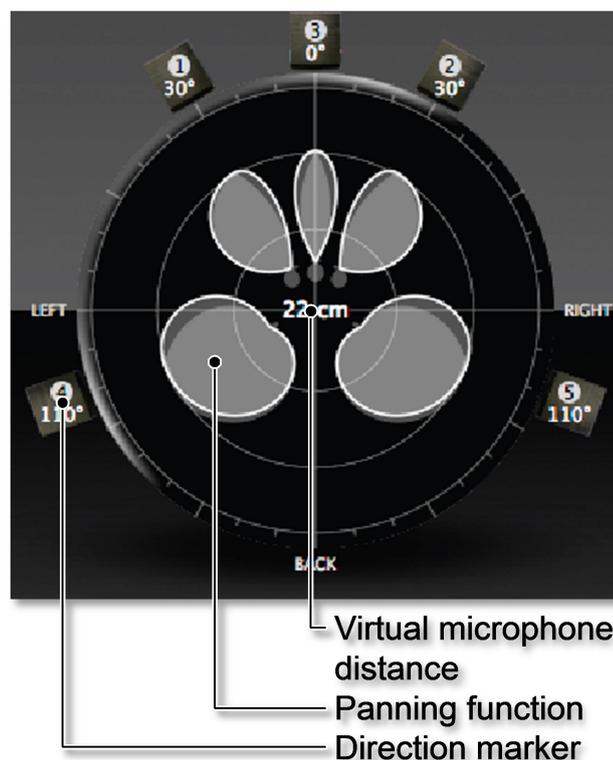
- **TwistedCube:** Eight loudspeakers arranged like the vertices of a twisted cube, i.e. the bottom loudspeakers are twisted 45° about the vertical axis.

1	2	3	4	5	6	7	8
Low front	Low left	Low rear	Low right	High front left	High rear left	High rear right	High front right

Surround

This output mode is suitable to decode a signal for playback over a horizontal loudspeaker array. Up to eight output channels can be activated with the +/- channel buttons. Each channel is associated with a loudspeaker. The direction of each loudspeaker can be set by clicking on its direction marker and dragging it to its new location.

The loudspeaker layout results in a set of panning functions. Different panning functions are



used in different frequency regions, which is reflected by the different shades of gray.

These panning functions may alternatively be thought of as the directivity patterns of tailor-made microphones. To add phase shifts corresponding to a spacing between these microphones, click on a panning function and drag it farther from the center of the diagram. The distance is displayed at the center.

An LFE channel may be selected by clicking the corresponding LFE button. The LFE dial can then be used to select the desired crossover frequency between the LFE channel and the other channels. To assign a silent LFE channel, turn the LFE dial to zero. Several LFE channels may be selected, in which case the low frequency content is panned between the LFE channels. Otherwise, the LFE channel will use an omnidirectional directivity function and the output is therefore independent of the location chosen for the LFE loudspeaker. The level of the LFE channel is not internally attenuated, but the surround presets will pull down the corresponding fader by 10 dB relative to other channels, in compliance with standards for surround sound distribution. Compliant playback equipment amplifies the LFE channel by 10 dB more than the other channels.

The surround decoder uses phase-matched shelf filters to apply slightly different panning laws to low frequencies and high frequencies. This does not apply to the LFE channel. Use the shelf dial to set the crossover frequency between the panning laws. For home listening, 400 Hz is recommended. For larger listening areas, the optimal crossover frequency is lower.

The following presets are available in this output mode:

- **Stereo, StereoSurround, StereoCenter, StereoSide:** Layouts with 2 loudspeakers. Used as an alternative to the stereo output mode.
- **3.0C/4.0C/5.0/6.0C/7.0C/8.0C3.1C/4.1C/5.1/6.1C/7.1C:** Output modes corresponding to “cinema” layouts, with loudspeakers more densely packed in the front.
- **5.0S/5.1S:** Surround layouts where the surround channels are directly at the sides.
- **3.0M/4.0M/6.0M/7.0M/8.0M**
3.1M/4.1M/6.1M/7.1M: Output modes corresponding to “music” layouts, with loudspeakers evenly spaced.
- **HexagonS, OctagonS:** Hexagonal and octagonal layouts with one speaker straight ahead.
- **HexagonF, OctagonF:** Hexagonal and octagonal layouts with the center of a speaker pair straight ahead.

AmbiX

This mode is used in case you want to use Harpex-X as an ambisonic processor, i.e. with ambisonic input and output and /or if you wish to upsample to higher-order ambisonics.

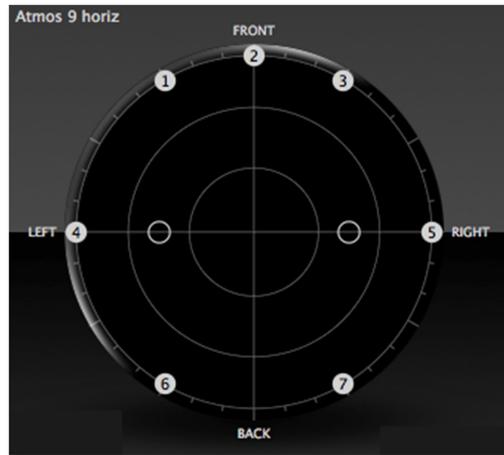
The first four output channels from Harpex-X will be the first-order signal. Discard the other twelve channels if you only want to use this signal. To access a

second order AmbiX signal, use the first nine channels and discard the last seven. Use all channels for a third order signal.

Speakerdome

This mode supports a number of standard 3D loudspeaker layouts by combining the Harpex method with VBAP / EBAP panning.

The supported loudspeaker layouts are listed below.



- **Octahedron:** Six loudspeakers arranged like the vertices of an octahedron.

1	2	3	4	5	6
Front left	Front right	Rear left	Rear right	Above	Below

- **3D7.0:** Seven loudspeakers forming a modified 7.0 layout.

1	2	3	4	5	6	7
High front left	High front right	Center front	High rear	Low front	Low rear left	Low rear right

- **Cube:** Eight loudspeakers arranged like the vertices of a cube.

1	2	3	4	5	6	7	8
Low front left	Low rear left	Low rear right	Low front right	High front left	High rear left	High rear right	High front right

- **TwistedCube:** Eight loudspeakers arranged like the vertices of a twisted cube, i.e. the bottom loudspeakers are twisted 45° about the vertical axis.

1	2	3	4	5	6	7	8
Low front	Low left	Low rear	Low right	High front left	High rear left	High rear right	High front right

- **Atmos 9:** Nine loudspeakers arranged in a 7.0 surround setup, supplemented with two overhead channels. Adapted to the Dolby Atmos® format.

Atmos 9								
1	2	3	4	5	6	7	8	9
Left	Center	Right	Left side	Right side	Left surround	Right surround	Above left	Above right

- **Auro 9:** Nine loudspeakers arranged in a 5.0 surround setup, supplemented with four overhead channels. Adapted to the 9 channel version of the Auro 3D® format.

Auro 9								
1	2	3	4	5	6	7	8	9
Left	Center	Right	Left surround	Right surround	High Left	High Right	High Left surround	High Right surround

- **Auro 10:** Ten loudspeakers arranged in a 5.0 surround setup, supplemented with five overhead channels. Adapted to the 10 channel version of the Auro 3D® format.

Auro 10									
1	2	3	4	5	6	7	8	9	10
Left	Center	Right	Left surround	Right surround	High Left	High Right	High Left surround	High Right surround	Above

- **Auro 11:** Eleven loudspeakers arranged in a 5.0 surround setup, supplemented with six overhead channels. Adapted to the 11 channel version of the Auro 3D® format.

Auro 11					
1	2	3	4	5	
Left	Center	Right	Left surround	Right surround	
6	7	8	9	10	11
High Left	High Center	High Right	High Left surround	High Right surround	Above

- **Auro 13:** Thirteen loudspeakers arranged in a 7.0 surround setup, supplemented with six overhead channels. Adapted to the 13 channel version of the Auro 3D® format.

Auro 13						
1	2	3	4	5	6	7
Left	Center	Right	Left side	Right side	Left surround	Right surround
8	9	10	11	12	13	
High Left	High Center	High Right	High Left surround	High Right surround	Above	

- **THX 10.0:** Ten loudspeakers arranged in a setup with five loudspeakers across the front half plane, two surround channels, one rear channel and two overhead channels.

THX 10.0							
1	2	3	4	5	6	7	8
30° left	Center	30° right	60° left	60° right	Left surround	Rear	Right surround
9	10						

Left above	Right above
------------	-------------

- **THX 12.0:** Twelve loudspeakers arranged in a setup with seven loudspeakers across the front half plane, two surround channels, one rear channel and two overhead channels.

THX 10.0							
1	2	3	4	5	6	7	8
30° left	Center	30° right	60° left	60° right	Left surround	Rear	Right surround
9	10	11	12				
Left above	Right above	90° left	90° right				

- **22.0:** Twenty-two loudspeakers arranged in a setup with ten loudspeakers at ear height, three below the screen and nine overhead. Requires two instances of Harpex-X running, one with the 22.0 A preset and the other with the 22.0 C preset.

22.0 A							
1	2	3	4	5	6	7	8
Left	Center	Right	Left surround	Right surround	Left side	Rear	Right side
9	10	11	12	13	14	15	16
Above left	Above right	Left side	Right side	Above front	Above left surround	Above right surround	Above rear
22.0 C							
1	2	3	4	5	6		
Above left side	Above	Above right side	Low left	Low right	Low center		

CPU usage

The real-time implementation of the Harpex method has only recently become a practical proposition, considering the large amount of computation which is required. On a modern computer, the CPU load of Harpex-X is usually a small fraction of the available resources, but may become prohibitive when used at high sampling frequencies and when using several instances in parallel. Certain measures can reduce the CPU load:

- All plugins handle blocks of samples. The size of these blocks can be configured in the DAW / soundcard driver. The larger these blocks are, the more evenly Harpex-X can distribute its load over time. In a real-time application like a DAW, this is essential to avoid drop-outs. Try to increase the buffer size / vector size / block size of your DAW / soundcard.
- Some of the most CPU intensive operations in Harpex-X are involved in the virtual microphone distance emulation. If this is not a required

feature, try to reduce the distance to zero (applies to stereo and surround output modes).

Notes pertaining to specific hosts

Pro Tools

If you receive error messages because you intermittently run out of CPU despite increasing the buffer size, try to check the box "Ignore Errors During Playback/Record" in the playback engine.

Reaper on Mac OS X

Although Reaper supports both VST and AU plugins, only the VST version of Harpex-X is compatible with Reaper on Mac OS X.

Audacity

Audacity currently does not support multi-channel plugins. Also, it seems that Audacity on OS X loads the Audio Unit version of Harpex every time it is started, which can be a nuisance during the trial period, when a dialog window is displayed at startup.

TwistedWave on Mac OS X

Although TwistedWave supports both VST and AU plugins, only the AU version of Harpex is compatible with TwistedWave on Mac OS X.

Plogue Bidule on Mac OS X

Although Plogue Bidule supports both VST and AU plugins, only the VST version of Harpex is compatible with Plogue Bidule on Mac OS X.

Logic Pro

To use Harpex in Logic Pro, create a surround track and insert Harpex on that track. To produce stereo output, go to the mixer and select "Output -> Stereo Output" under the "I/O" header of the track where Harpex is inserted.

Nuendo

When placing Harpex on an ambisonic track, it is necessary to select the ambisonic input format in the upper right-hand corner of the plugin window, even though this seems to be selected by default. Also note that the ambisonic channels, named 0-16, are not by default routed to loudspeaker outputs. For this, a new bus needs to be created in ambisonic format where each channel is assigned to an output.