

Wave Arts Panorama 7



User Manual

Last updated: December 2, 2022
Copyright © 2022, Wave Arts, Inc. All Rights Reserved

Table of Contents

1.0 Introduction	3
1.1 Overview of interface	4
1.2 Quickstart.....	6
1.3 Uses of Panorama 7	7
1.4 What's new in Panorama 7	8
2. Using Panorama.....	9
2.1 Main section.....	9
2.2 Main display.....	12
2.3 Direct/Slot controls	13
2.4 Reflections.....	14
2.5 Reverb	14
2.6 Spatialization	16
2.7 Sequencer	17
3. Theory and implementation	20
3.1 About 3D Audio and Acoustic Environment Modeling	20
3.2 Panorama Audio Routing	29
3.3 Parameters	31
4. Installation and Registration	34
4.1 Installation	34
4.2 Registration	34
4.3 Wave Arts licensing	34
4.4 PACE/iLok licensing	34
5. Plug-in Control Operation	37
5.1 Knobs	37
5.2 Text Entry	37
5.3 Sliders	38
6. Menu Bar and Preset Manager	39
6.1 Enable	39
6.2 Undo/Redo	39
6.3 Preset name and arrow controls	39
6.4 Preset menu	39
6.5 A/B buffers	42
6.6 Copy buffer.....	42
6.7 Tools menu.....	43
Trademark Notices.....	45
License Agreement.....	46
Support.....	49
Index.....	50

1.0 Introduction

Panorama is a virtual acoustics processor that can produce stunningly realistic auditory scenes for playback over headphones or loudspeakers. Panorama uses 3-D audio and acoustic environment modeling technology to simulate the acoustics of sound propagation in real spaces.

Panorama features the following:

- Head-related impulse response based binaural synthesis
- Early room reflection modeling
- Late diffuse reverberation
- Doppler pitch effect
- Distance modeling
- Position sequencer for automating trajectories and creating musical panning effects (auto-pan)
- Crosstalk canceling based on real head models for loudspeaker playback
- Support for SOFA format HRIRs
- Support for user supplied HRIRs
- Conventional stereo panning mode

1.1 Overview of interface



1. Menu bar. Bypass, Undo/Redo, Presets, A/B buffers, and Settings
2. Main display. Move the cursor (green ball) to move the sound. Select top, side, rear, or 3D display, and zoom factor.
3. Spatializer or Sequencer display selection.
4. Output configuration selection. **Very important to set this first.**
5. Head-Related Impulse Response selection.
6. Sequencer enable. Click this to start/stop sequencer playback.
7. Direct / Slot parameters. Shows the current selected cursor or waypoint coordinate.
8. Reflections parameters.
9. Reverb parameters.

10. Spatialization section. Mix the direct sound, reflections, and reverb.



11. Sequencer section. Set up waypoints to automate the sound trajectory.

1.2 Quickstart

It's very important to set the "OUTPUT" configuration to match your listening setup. So, if using headphones, make sure the OUTPUT is set to Headphones, or, if listening over loudspeakers, set one of the loudspeaker configurations. The loudspeaker width is an angle measured from the center, hence the 30 deg setting corresponds to a typical stereo setup where the listener and the two speakers form an equilateral triangle. Use the 10 degree setting for closely spaced speakers, say on a laptop computer.

The default preset "Hdphs – mono, anechoic" sets headphone configuration, mixes stereo inputs to mono, and disables reflections and reverb. It's a good place to start. We recommend listening over headphones the first time you use Panorama.

Play some music or speech and then drag the cursor (green ball) in the main display. As you move the cursor, you should hear the sound move around your head in a realistic way that is not possible with regular stereo panning.

Assuming the main display is showing the "TOP" view, moving the cursor will move the sound around in the horizontal plane. Use your mousewheel to move the sound up and down. You can also directly manipulate the x, y, and z controls (if in Cartesian mode) or the azimuth, elevation, and distance controls (if in Polar mode).

You will note the sound gets louder when close to the head which provides a distance cue. Now enable Reflections and you will hear spatialized early reflections off the virtual walls and floor and ceiling – this gives a real sense of space. Now enable Reverb and hear the diffuse late reverberation. You can mix the relative levels of the direct, reflections, and reverb in the Spatialization section (using the GAIN knobs), and also adjust how the loudness of each component changes as a function of distance (using the DIST ROLLOFF knobs).

If your input sound is stereo, increase the WIDTH control in the DIRECT section. The stereo sound will be spatialized using two virtual loudspeakers positioned on each side of the cursor.

You can also enable the DOPPLER effect and try moving the sound past the head. If your input sound is tonal you should hear the Doppler pitch effect as the moving sound passes your head.

The Sequencer can be shown via the SEQUENCER/SPATIALIZER selection, clicking SEQUENCER shows the Sequencer UI. Try selecting the "Orbit Clockwise" preset and then click the "SEQUENCER" toggle button to start playback. This will automate the cursor position in a circular arc from the user's left around the front to the right and then around the rear of the

listener. Note that the sequencer operates only when audio is streaming through the plug-in, and the SEQUENCER toggle is enabled.

1.3 Uses of Panorama 7

There are many possible uses for Panorama.

- **Virtual scene synthesis.** You can insert Panorama on each track in a mix, set up the same reflection and reverb settings in each instance, and then pan the individual sounds around in the virtual space. This works better for headphone mixes (think of the possibilities for podcasting), but the same concept can be used for loudspeaker playback. You can use the sequencer to set up motion trajectories for the sources, or simply automate them.
- **Stereo panning effects.** In conventional stereo mixes, sometimes you have a sound you want to treat specially. Panorama can be used to create a panning effect that stands out from conventional power panning. You can for example select a loudspeaker mode to enable the crosstalk canceller and then position a sound to an extreme lateral position like +/- 90 degrees. The sound will appear outside the regular stereo soundfield. You can also use headphone mode and simply apply binaural cues to the sound – even over loudspeakers the binaural spatial effect is oddly different than regular stereo panning.
- **Stereo widening.** Panorama can process stereo mixes to sound very wide when played back over loudspeakers, and this can work very well when the playback speakers are close together such as on a laptop. To do this, insert Panorama on the stereo master bus, set the output to loudspeaker 10 degrees, set the source to front-center and set the width to 90 degrees. There are presets for this.
- **Stereo “auto-pan” effects.** Panorama can work as a conventional stereo panner (select the Stereo Pan HRIR), and it has a powerful sequencer to automate pan positions. Hence you can create all sorts of modulating pan effects. You can enable Doppler to include pitch modulation. Adding reflections can sound like a choruser.
- **Virtual stereo monitoring.** To hear simulated stereo loudspeaker playback over headphones, insert Panorama on the master bus, position the cursor at front center, and select the width of your virtual speakers, say 30 deg. There is a preset for this.
- **Virtual surround monitoring.** Using the same idea you can simulate surround speaker setups, such as 5.1, 7.1, or 7.4.1. Because Panorama doesn’t support surround inputs, you have to split the surround input into L, R stereo pairs and insert Panorama on each pair. Center and subwoofer can be combined on a mono channel rendered using a front-center source. The 5.1 surround channels would be positioned at front-center with a stereo width of 110 deg, so the virtual surround speakers will be at +/- 110 deg. Repeat the

positioning for each speaker pair, adding elevation as needed. This works for both headphone and stereo loudspeaker playback.

- **Binaural to loudspeaker conversion.** A binaural recording is made using tiny microphones placed in the ear canals of a listener. To playback such recordings over loudspeakers you want to process them with a crosstalk canceller. Panorama has presets for doing this which uses the Stereo Passthru HRIR to bypass binaural synthesis.


1.4 What's new in Panorama 7

- The big new feature is the position sequencer, which can be used for trajectory automation and beat-synced musical effects (auto-pan). The sequencer works by panning the sound along a set of trajectory waypoints.
- All new UI look with dark theme.
- Larger UI area for panning, now including side and rear view options.
- Larger virtual area for panning. Previously Panorama 6 allowed from -10 to +10 feet, Panorama 7 allows from -100 to +100 feet.
- Zoom in/out control.
- Support for SOFA format HRIRs. SOFA (Spatially Oriented Format for Acoustics) is a standard file format used for interchange of head-related impulse responses. There are large publicly available sets of binaural responses in SOFA format. Support for SOFA data is provided via a conversion application called "Sofa2Pan" that converts from SOFA format to Panorama format.
- Improved support for third party HRIRs, HRIRs are now specified by filename rather than integer ID.
- HRIRs are now sample rate converted to the session rate when loaded. Previously they had to be provided at the proper sampling rate.
- Implemented "Near Head Model". This improves reproduction accuracy for positions close to head by applying additional head shadowing based on a spherical head model.
- Tweaked the reverb algorithm to remove early direction bias.
- Added "Stereo Pan" HRIR to implement traditional stereo panning.
- Added "Stereo Passthru" HRIR to implement proper binaural to loudspeaker conversion.
- Switching between Feet and Meters units preserves positions and distances, previously using metric units also scaled the distances.
- Internal coordinate parameters are all cartesian, previously both polar and cartesian parameters were exposed along with a cartesian/polar toggle.
- Improved preset management with infinite undo/redo which includes view options.

2. Using Panorama

This chapter will describe the Panorama user-interface, assuming the user is familiar with the technology behind spatial/3D audio. If you are not yet familiar, please skip ahead to the next chapter which gives a detailed explanation of the theory of 3D audio and the implementation of Panorama.

2.1 Main section



The SPATIALIZER/SEQUENCER button is used to change the UI view between the spatializer controls and sequencer controls.

The OUTPUT control sets the playback configuration. It's very important to set this correctly. The current setting is depicted in the graphic display. When set to loudspeaker playback, Panorama will process the output signal with a "crosstalk canceller" tailored for the specified speaker width. The width is the angle from the front center position to each speaker. The possible output configurations are:

- **Headphones** – use for headphone playback, regardless of headphone type. Crosstalk canceller is disabled.
- **Speakers, 10 degrees** – use for playback over closely spaced speakers, such as found on laptop computers. This setting has the most exaggerated effect of the speaker settings.
- **Speakers, 20 degrees** – use for playback over desktop computer speakers.
- **Speakers, 30 degrees** – use for playback over standard stereo setups.
- **Speakers, 40 degrees** – use for playback over stereo setups with widely spaced speakers. This setting has the least exaggerated effect of the speaker settings.

Unfortunately, the output modes are not entirely compatible with both headphones and loudspeakers. Panorama is ideally applied when you know what the playback configuration will be. What configuration should be used if you are not sure?

If Panorama is being used to create a special track effect (e.g., a flying sound) for a standard stereo production, we recommend using the speakers

setting, 20 or 30 degrees, because the result will probably be listened to over speakers. The flying track will stand out from the rest of the stereo mix created using standard stereo panning.

If you want to create a special effect for headphone listeners (e.g. iPhone listeners), use the headphone mode. This mode gives the most realistic 3D experience.

Listening to headphone mode audio over loudspeakers is not bad, but you won't get the really wide sounds that are possible using crosstalk cancellers. In contrast, listening to crosstalk cancelled audio over headphones can sound pretty unnatural. Crosstalk cancellation creates out-of-phase low frequencies which can impart an unusual pressure sensation to the ears when listened to over headphones. This effect is more noticeable if the entire mix has been processed with a crosstalk canceller.

The HRIR control lets the user select the head-related impulse response set. The popup menu has the following options:

- **Human.** This is a set of HRIRs measured on a human subject, specifically subject 003 from the CIPIC data set. Wave Arts conducted localization experiments using a number of listeners and HRIRs and this HRIR set was found to perform best in terms of front/back localization perception.
- **Human2.** This is the HRIR set of Panorama's author.
- **Kemar.** This is a measurement of a KEMAR (Knowles Electronic Manikin for Acoustic Research) measured by the author at the MIT Media Lab. While it's a well-known set of data, it's not particularly good for localization, in particular it has very poor frontal localization perception. It is included for compatibility with earlier Panorama software. It's also potentially useful to impart rear localization perception when using loudspeakers.
- **Stereo Pan.** This is a precomputed HRIR which implements conventional stereo panning without filtering or interaural time delay. The pan position is based on the azimuth: -90 is left, -30 is left-center, 0 is center, +30 is right-center, +90 is right. The pan function is equal power, i.e., cosine and sine curves with pan gains of -3 dB at center.
- **Stereo Passthru.** This is a precomputed HRIR which bypasses the binaural synthesis processing by passing any left hemisphere source to left output without filtering or delay, and does the same for righthand sources. It is used for the presets that implement binaural to loudspeaker conversion, where the input signal already contains binaural cues. The source azimuth should be at least +/- 5 degrees, because center positions are muted.
- **Factory.** This is a rolloff menu listing all HRIRs found in the factory HRIR directory. Panorama currently ships with the complete set of CIPIC HRIRs installed in a "CIPIC" sub-directory. On MacOS the factory directory is "/Library/Application Support/Wave Arts/Panorama

- 7/PanoramaHRTFs". On Windows the factory directory is "C:\ProgramData\Wave Arts\Panorama 7\PanoramaHRTFs".
- **User.** This is a rolloff menu listing all HRIRs found in the user HRIR directory, which is initially empty. By default the SOFA conversion will leave HRIRs in this directory. On MacOS the user HRIR directory is "~/Library/Application Support/Wave Arts/PanoramaHRTFs" where "~/Library" is your user library directory. On Windows the user HRIR directory is "C:\Users\USER\AppData\Roaming\Wave Arts\PanoramaHRTFs" where USER is your user login name.
 - **Import SOFA data.** This will launch the Sofa2Pan application if it is found on your computer. Please see the documentation that accompanies Sofa2Pan.

The CART/POLAR button selects cartesian or polar coordinate display in the UI. In cartesian mode, coordinates are specified in terms of x, y, and z. In polar mode, coordinates are specified in terms of azimuth, elevation, and distance. Note that internally, Panorama 7 parameters are all cartesian, hence when automating parameters you will see only cartesian parameters listed.

The SEQUENCER toggle starts/stops sequencer playback. The toggle button also appears in the Sequencer section. It is provided in the main section so that you can start/stop sequencer playback when viewing the spatializer parameters.

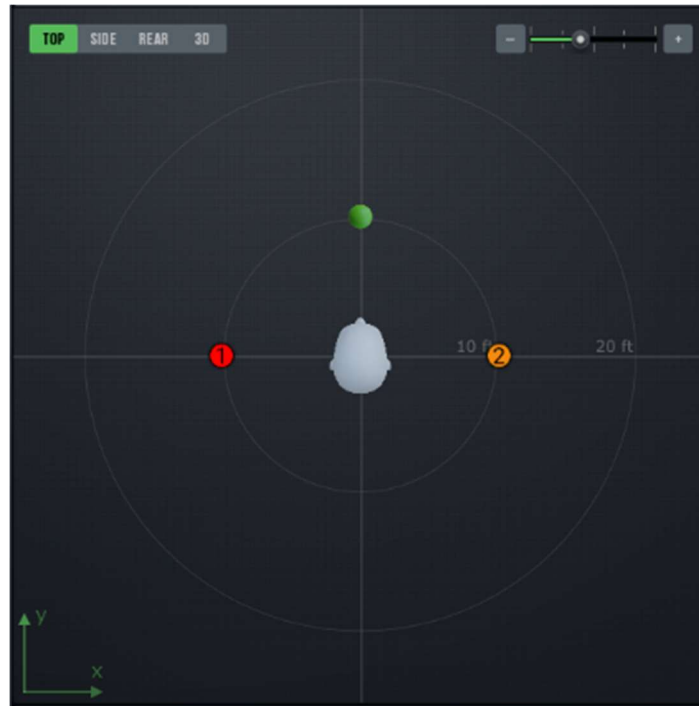
The DOPPLER toggle enables/disables the Doppler pitch effect. When Doppler is enabled, the sound propagation from each virtual source to the listener is simulated using a variable delay line where the delay time is the distance divided by the speed of sound; this naturally reproduces the Doppler pitch effect. When Doppler is disabled, sound propagation delay is simulated as if the source is at the listening position (origin), while the sound propagation gain is simulated normally. Hence, with Doppler disabled, the direct source has a delay of 0 and the room reflections have delays corresponding to twice the wall distances, regardless of the position of the source.

The NEAR HEAD MODEL toggle enables/disables the near head model. Near head model increases head shadowing when the source distance is closer than 1.5 meters, it has no effect at source distances greater than 1.5 meters.

The GAIN control applies gain to the Panorama output.

The output meters show output levels along with a peak hold indicator.

2.2 Main display



The main display depicts the virtual space where you pan the input sound. The “cursor” is the green ball which corresponds to the current 3D pan position. You can click and drag this to pan the sound around the space.

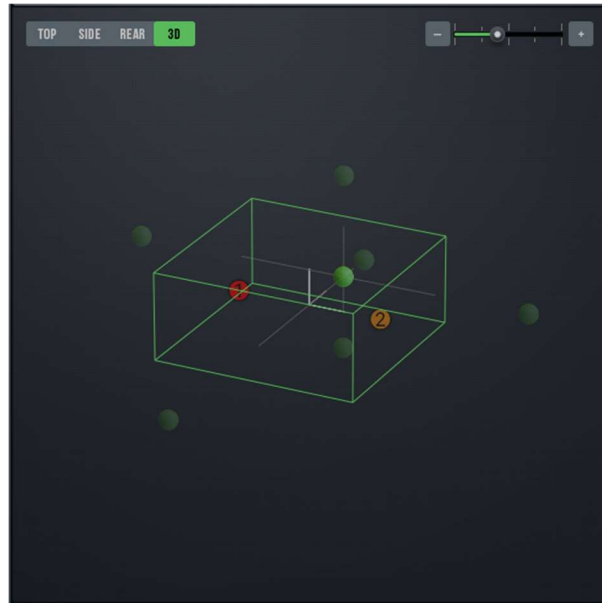
If the sequencer is enabled or the sequencer UI is visible, the main display will show the sequencer “waypoints” depicted as colored balls with numbers. During sequencer playback, the cursor moves along a trajectory defined by the waypoints.

The TOP/SIDE/REAR/3D buttons select the current viewpoint. Top is an overhead view, Side is from the right side, Rear is from behind. These are all orthographic projections (no perspective) although there is a perspective effect applied to the size of the controls – they increase in size when closer to the viewer. The head is drawn at maximum zoom factor and stays at this size even as the zoom is decreased.

The zoom control lets you zoom in/out in factors of 2. When fully zoomed out one can see the entire +/- 100ft range, when zoomed in the range is about +/- 3 feet. If the zoom causes a control point to go offscreen, an arrow indicator is drawn at the edge pointing to the offscreen control.

The 3D setting shows a 3D depiction of the space, along with a wireframe drawing of the current room dimensions. The origin is shown with white axis lines along the x, y, and z axes. The green lines represent the current room. The faint green balls are the current virtual reflection locations. Clicking and

dragging in the 3D room changes the viewpoint – there is no way to drag the control points in the 3D view.



2.3 Direct/Slot controls



The Direct controls change the pan position using x, y, z, coordinates when in Cartesian mode, and elev, azimuth, and dist coordinates when in Polar mode. When editing sequencer waypoints, the section label will read "Slot N" where N is the number of the slot being edited.

The Direct toggle control enables/disables the direct sound, i.e. the propagation from the source directly to the listener. When disabled, the direct source is muted, allowing the user to audition just the early reflections or reverberation.

The Direct section contains a preset control used to save and recall pan positions. The factory presets include the standard 7.4.x surround locations. Note that the Width control is neither saved nor recalled in the Direct presets.

2.4 Reflections



The Reflections control lets you change the dimensions of the virtual room, select wall surface materials, and enable/disable individual reflections.

The Reflections toggle enables/disables the room reflections in the output mix.

The preset control lets you save and recall room presets.

Note that if the cursor position moves outside a wall surface, the corresponding virtual reflection will move into the room. Furthermore, there is no occlusion effect applied to the direct source. This is not physically accurate but was the most practical approach, hence it is up to the user to restrict cursor position to the interior of the virtual room.

2.5 Reverb



The Reverb control lets you change the late reverb parameters in the virtual room. The reverb is parameterized using decay time, room size, pre-delay, and early and late damping:

Reverb Decay Time — Sets the decay time of the reverberation in seconds, from 0.5 sec to 60 sec. Also known as the reverberation time, the decay time is the time it takes for frequencies near 500 Hz to decay by 60 dB. For typical input sounds, the reverberation will be essentially inaudible after decaying 60 dB, so the Decay Time specifies the effective length of the reverberation tail in seconds. The decay time parameter is independent of the room size.

Reverb Room Size — Scales the size of the simulated room, ranging from 0% to 100%. A room size of 0% sounds like a very small space, such as a bathroom or closet, 25% sounds like a small studio, 50% sounds like a jazz club, 75% sounds like a concert hall, and 100% sounds like a very large room, such as an aircraft hangar or stadium. The room size parameter is independent of the decay time. In general, small rooms sound best with short decay times, and large rooms sound best with longer decay times.

Reverb Pre-Delay — Delays the reverb output. Depending on the sound source, small delays, less than 50 msec, will generally increase the audibility of the reverberation. Larger delays, greater than 80 msec, will usually be perceived as a separate echo. Hence the pre-delay parameter can be used to create an echo effect for simulating really big spaces.

Reverb Early Damping — Adjusts the cutoff frequency of a lowpass filter at the input of the reverb; thus, the Early Damping parameter sets the upper limit of frequencies admitted into the reverb. The range of the Early Damping parameter is from 20 Hz to 20 kHz. Setting this parameter to the maximum value opens up the filter completely for maximum brightness. A low Early Damping setting results in a muffled reverberation, whereas a high Early Damping setting results in a bright, crisp reverberation.

Reverb Late Damping — Sets the cutoff frequency of the damping filter in the reverb. A lower setting causes high frequencies to decay faster, resulting in a warmer sounding late reverberation; a higher setting results in a brighter sounding late reverberation. The range of the Late Damping is from 20 Hz to 20 kHz. Setting this parameter to the maximum value opens up the filter completely for maximum brightness. A value of around 4 kHz results in a natural sounding reverberation for typical rooms. Rooms with hard reflective walls consisting of concrete or tile, such as a bathroom, would be simulated using a higher Late Damping. Rooms containing much absorbent material would be simulated using a lower Late Damping.

The Reverb toggle enables/disables the reverb in the output mix.

The preset control lets you save and recall reverb presets.

Note that the reverb and reflections settings are completely independent. Typically, one would use similar sounding setups, e.g., a short reverb time with small size would ideally be paired with room reflections from a similarly small space.

2.6 Spatialization



The spatialization control can be thought of as a mixer which combines the direct sound, reflections, and late reverb to form the output sound. Each of the three components has an enable toggle, which mirrors the toggles found in the sections, a gain control which sets the mix level, and a distance rolloff control. The distance rolloff specifies the attenuation in sound level per doubling of distance from the listener, starting from a reference point of 1 foot, where the level is specified by the corresponding gain parameter. A value of 0 dB will cause the source to have the same gain regardless of distance.

For the direct and reflections sources, a rolloff value of -6 dB simulates natural acoustics, but may seem exaggerated, hence the Panorama presets use -4 dB.

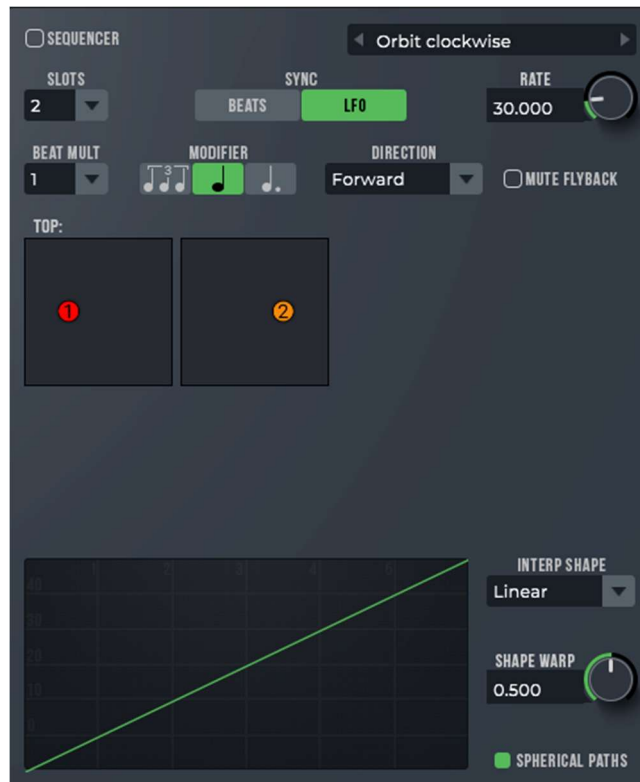
The reverb level stays relatively constant in rooms as a function of distance, decreasing slightly due to absorption. The Panorama presets use -1.5 dB.

Chapter 3 has additional discussion of distance perception.

The HRIR length controls specify the maximum filter lengths (in taps) for the direct and reflection virtual sources. Longer filter lengths are more accurate, reproducing finer spectral detail, but take more CPU power.

We recommend using shorter filters for the reflections than for the direct source, for two reasons: 1) the reflections are not as important to reproduce accurately, and 2) there are up to six reflections, so the reflections can consume six times the CPU as the direct source when the filter lengths are the same.

2.7 Sequencer



The sequencer lets you automate the pan position by setting up a set of waypoints that define the trajectory of the cursor. Each waypoint is associated with a sequencer timeslot, which corresponds to a musical timespan such as a beat or a measure. You first set the number of timeslots desired, then you set the position of each slot's waypoint. Enabling the sequencer starts playback; the pan position will move along the set of waypoints and will then repeat according the direction parameter.

The SLOTS parameter sets the number of timeslots, from 2 to 8. Newly created slots are shown as "empty" which correspond to musical rests. During playback, trajectories interpolate through empty slots to the next available waypoint. Click on a slot to select it, then click again to create a waypoint in the slot. You can then position the waypoint in the slot, or for finer control, drag it in the main view, or use the position knobs. Typing delete will delete the waypoint in the current slot, leaving it empty. The waypoint in slot 1 cannot be deleted because sequences need at least one waypoint.

The sequencer is always synchronized to the host DAW transport. The SYNC source is either BEATS or LFO; BEATS uses the host's musical transport position, LFO generates beats based on a rate parameter in beats-per-minute. In either case, the sequencer **will not play** unless the host's transport is running and actively processing samples.

The BEAT MULT control lets you speed up or slow down playback by factors of two. When the multiplier is 1, the sequencer advances one timeslot per beat, a multiplier of 4 means that time advances one slot per 4 beats. The BEAT MODIFIER control selects triplets, beats, or dotted notes, hence contributing a factor of $x2/3$ for triplets, 1 for beats, and $x3/2$ for dotted notes.

The DIRECTION parameter controls the direction and repeat mode, and has the following options:

- **Forward.** Sequence advances from slot 1 to N and then repeats at 1.
- **FwdJump.** Sequence advances from slot 1 to N but immediately jumps back to slot 1 as the transport reaches slot N.
- **Backward.** Sequence moves backwards from slot N to 1 and then repeats at slot N.
- **BkwdJump.** Sequence moves backwards from slot N to 1 but immediately jumps back to slot N as the transport reaches slot 1.
- **FwdBkwd.** Sequence advances from slot 1 to N but changes direction as the transport reaches slot N, then moves backward to slot 1 and changes direction again at the start of slot 1.

The MUTE FLYBACK toggle causes the sound to mute during the “flyback” portion of the sequence, which is between slot N and 1, assuming slot N and slot 1 both have waypoints. In case either slot is empty, the flyback is between the nearest waypoints that bound the slot N to slot 1 transition. The FwdBkwd mode has no flyback portion. The Jump modes have a very brief flyback (20 msec).

The slot displays share the same viewpoint and zoom as the main view; the viewpoint is shown above slot 1. You can independently select the slot viewpoint by CTRL-clicking on the TOP/SIDE/REAR buttons. This is useful when the main view is showing the 3D room.

Normally the trajectory moves along the waypoints in a linear manner. The INTERP SHAPE and SHAPE WARP parameters let you modify this. The INTERP SHAPE has the following values:

- **Linear.** Perfect linear interpolation when the warp is 0.5. For warps less than 0.5, the trajectory waits at the source slot for a while before moving. For warps greater than 0.5, the trajectory moves to the destination and then waits.
- **Exponential.** Perfect linear interpolation when warp is 0.5. For warps less than 0.5 the trajectory slowly moves away from source. For warps greater than 0.5, the trajectory quickly moves away from source.
- **Sinusoidal.** A perfect sinusoidal shape when warp is 0.5. This interpolates to a linear shape at warp of 1.0, and to a linear segment shape at warp of 0.

The SPHERICAL PATHS toggle, when enabled, causes the trajectories to move along circular arcs instead of straight lines. Essentially, the inter-waypoint interpolation is done using polar coordinates rather than cartesian.

3. Theory and implementation

The following sections provide a tutorial on 3-D audio and acoustic environment modeling, followed by a description of Panorama's implementation.

3.1 About 3D Audio and Acoustic Environment Modeling

3.1.1 What is 3D Audio?

A 3D audio system has the ability to position sounds all around a listener. The sounds are actually created by the loudspeakers (or headphones), but the listener's perception is that the sounds come from arbitrary points in space. This is similar to stereo panning in conventional stereo systems: sounds can be panned to locations between the two loudspeakers, creating virtual or "phantom" images of the sound where there is no loudspeaker. However, conventional stereo systems generally cannot position sounds to the sides or rear of the listener, nor above or below the listener. A 3D audio system attempts to do just that.

To understand how 3D audio systems work, it is useful to start by considering how humans can localize sounds using only two ears. A sound generated in space creates a sound wave that propagates to the ears of the listener. When the sound is to the left of the listener, the sound reaches the left ear before the right ear, and thus the right ear signal is delayed with respect to the left ear signal. In addition, the right ear signal will be attenuated because of "shadowing" by the head. Both ear signals are also subject to a complicated filtering process caused by acoustical interaction with the torso, head, and in particular, the pinna (external ear). The various folds in the pinna modify the frequency content of the signals, reinforcing some frequencies and attenuating others, in a manner that depends on the direction of the incident sound. Thus an ear acts like a complicated tone control that is direction dependent. We unconsciously use the time delay, amplitude difference, and tonal information at each ear to determine the location of the sound. These indicators are called sound localization "cues."

3.1.2 Head-Related Impulse Responses (HRIRs)

The transformation of sound from a point in space to the ear canal can be measured accurately; the measurements are called Head-Related Transfer Functions (HRTFs). The corresponding time domain responses are called Head Related Impulse Responses (HRIRs). The two acronyms are interchangeable. Because Panorama uses the time domain data, we use HRIR to refer to head response data in this document.

HRIR measurements are usually made by inserting miniature microphones into the ear canals of a human subject or a manikin. A measurement signal is played by a loudspeaker and recorded by the microphones. The recorded signals are then processed by a computer to derive a pair of HRIRs (for the left and right ears) corresponding to the sound source location. This process is diagrammed in the figure below.




Figure 3.1. HRIR measurement for a single source location.

Each HRIR, typically consisting of several hundred numbers, describes the time delay, amplitude, and tonal transformation for the particular sound source location to the left or right ear of the subject. The measurement procedure is repeated for many locations of the sound source relative to the head, resulting in a database of hundreds of HRIRs that describe the sound transformation characteristics of a particular head.

3.1.3 Binaural Synthesis

A 3D audio system works by mimicking the process of natural hearing, essentially reproducing the sound localization cues at the ears of the listener. This is most easily done by using a pair of measured head-related impulse responses (HRIRs) as a specification for a pair of digital audio filters (equalizers). When a sound signal is processed by the digital filters and listened to over headphones, the sound localization cues for each ear are reproduced, and the listener should perceive the sound at the location specified by the HRIRs. This process is called binaural synthesis (binaural signals are defined as the signals at the ears of a listener).




Figure 3.2. Binaural synthesis of a single source.

Binaural synthesis works extremely well when the listener's own HRIRs are used to synthesize the localization cues. However, measuring HRIRs is a complicated procedure, so 3D audio systems typically use a single set of HRIRs previously measured from a particular human or manikin subject.

3.1.4 Non-Individualized HRIRs

Localization performance generally suffers when a listener listens to directional cues synthesized from HRIRs measured from a different head, called "non-individualized" HRIRs. Human heads are all different sizes and shapes, and there is also great variation in the size and shape of individual pinnae. This means that every individual has a different set of directional cues. The greatest differences are in the tonal transformations at high frequencies caused by the pinnae. It is clear we become accustomed to localizing with our own ears, and thus our localization abilities are diminished when listening through another person's ears. Our uniqueness as individuals is a significant limitation of practical 3D technology.

The use of non-individualized HRIRs results in two particular kinds of localization errors commonly seen with 3D audio systems: front/back confusions and elevation error. A front/back confusion results when the listener perceives the sound to be in the front when it should be in back, and vice-versa. An elevation error refers to a misperceived elevation angle, for example an overhead source may be perceived as being in front of or behind the listener. Both front/back and elevation performance is much better when using headphones than when using loudspeakers because the high frequency cues are more faithfully reproduced.

3.1.5 Panorama HRIRs

The HRIRs used by Panorama were measured from human subjects and a dummy head microphone. The human measurements were made at the CIPIC Interface Laboratory at the University of California at Davis. Panorama's built-in "Human" set is derived from CIPIC Subject 003. This was experimentally determined to be a good non-individualized set for front/back resolution. The "Human2" set is from Panorama's author. The dummy head measured was a Knowles Electronic Manikin for Acoustic Research (KEMAR). The KEMAR was measured at the MIT Media Laboratory.

A variety of measurements are supplied with Panorama to allow users to audition different head models. It's possible that users of Panorama will find a particular HRIR set works better for them than the others. Which HRIR set should be used in production? There's no easy answer to this question. We have found the "Human" set works well for headphone reproduction, and the "KEMAR" set works well for loudspeaker reproduction.

3.1.6 Reproducing 3D Audio Using Loudspeakers

When reproducing localization cues to a listener, it is important that the left and right audio channels remain separated, that is, the left ear signal should go to the listener's left ear only, and the right ear signal should go to the listener's right ear only. This is easy to achieve when the listener is using headphones. When using loudspeakers, however, there is significant "crosstalk" between each speaker and the opposite ear of the listener. A large portion of the left speaker signal will go to the right ear of the listener, and similarly a large portion of the right speaker signal will go to the left ear of the listener. Crosstalk is depicted in the figure below as the acoustic paths from each speaker to the opposite ear. The crosstalk severely degrades localization performance and must be eliminated.




Figure 3.3. Crosstalk when listening to stereo loudspeakers.

3.1.7 Crosstalk Cancellation

Fortunately, it is possible to build a digital filter, called a “crosstalk canceller,” that eliminates crosstalk. The crosstalk canceller adds a cancellation signal to each of the two channels of audio, such that when the listener is properly positioned between the loudspeakers, the crosstalk is acoustically cancelled at the listener’s ears. The listener must be centered between the two loudspeakers in order for the crosstalk to be cancelled. In 3D audio parlance, the listener must be in the “sweet spot” to get the full 3D effect. Provided the listener is centered between the loudspeakers, crosstalk cancellation is relatively insensitive to front-back motions of the listener, however, crosstalk cancellation is degraded when the listener is off-center or not facing forward.

Loudspeaker 3D audio systems are extremely effective in desktop computing environments. This is because there is usually only a single listener (the computer user) who is almost always centered between the speakers and facing forward towards the monitor. Thus, the primary user gets the full 3D effect because the crosstalk is properly cancelled. In typical 3D audio applications, like video gaming, friends may gather around to watch. In this case, the best 3D audio effects are heard by others when they are also centered with respect to the loudspeakers. Off-center listeners may not get the full effect, but they still hear a high quality stereo program with some spatial enhancements.

Many crosstalk cancellers are based on a highly simplified model of crosstalk, for example modeling crosstalk as a simple delay and attenuation process, or a delay and a lowpass filter. Other crosstalk cancellers have been based on a spherical head model. The crosstalk canceller used by Panorama is based on actual HRIR measurements and thus accurately models the crosstalk that occurs with human listeners. For typical human listeners, the Panorama crosstalk canceller improves channel separation by about 20 dB in the 100 Hz to 6 kHz range. This may seem like a modest improvement, but in fact it is quite good. Even a small improvement in channel separation leads to a large improvement in localization performance. As with binaural synthesis, crosstalk cancellation performance is ultimately limited by the variation in the size and shape of human heads.

Because of the limitations of crosstalk cancellers, 3D audio over loudspeakers is prone to both front/back confusions and elevation errors. In particular, back to front confusions tend to be very common, which means that many listeners may not be able to perceive sounds as being in the rear. In practice, this means that when panning a sound from the front, around to the side, and to the rear, the result will be perceived as a sound panning to the side and then back to the front. Also, elevation performance is generally not good over loudspeakers. For example, when a sound is moved from directly to the right to directly overhead, this may be perceived as though the sound is moving from the right to directly in front.

3.1.8 Acoustic Environment Modeling

Acoustic environment modeling refers to combining 3D spatial location cues with distance, motion, and ambience cues, to create a complete simulation of an acoustic scene. By simulating the acoustical interactions that occur in the natural world, we can achieve stunningly realistic recreations, above and beyond that possible with just 3D positional control. Panorama combines 3-D positioning with accurate simulations of room surface reflections, late diffuse reverberation, distance cues, and the Doppler effect. These phenomena are described in the following sections.

3.1.9 Room Reflections and Reverberation

When an object in a room produces a sound, a soundwave expands outward from the source reaching walls and other objects where sound energy is both absorbed and reflected. Technically speaking, all reflected energy is called reverberation. Assuming a direct path exists between the source and the listener, the listener will first hear the direct sound, followed by reflections off nearby surfaces, called early reflections. After a few tenths of a second, the number of reflected waves becomes very large, and the resulting reverberation is characterized by a dense collection of soundwaves traveling in all directions, called diffuse reverberation. The time required for the reverberation to decay 60 dB below the initial level is defined as the reverberation time. Generally, reverberation in a small room decays much faster than reverberation in a large room, because in a small room the soundwaves collide with walls much more frequently, and thus are absorbed more quickly, than in a large room.

Reverberation is an important acoustic phenomena. There is at most one direct path from the source to the listener, whereas there may be millions of indirect paths, particularly in a room where a sound can bounce around hundreds of times before being absorbed. Thus, in typical listening situations, most of the energy we hear from a sound source is actually reflected energy.

The perception of reverberation depends on the type of reverberation and the type of sound. In small room with fast decaying reverberation, the reverberation imparts a tonal quality to the sound that is readily identified as a small room signature. In a larger room, the reverberation can create a background ambience that is easily distinguished from the foreground sound, and this is readily identified as a characteristic of large spaces. In this manner, reverberation imparts useful spatial information about the size of the surrounding space.

Reverberation that contains a lot of high frequency energy in the decay is associated with rooms that have hard, reflective walls, which do not readily absorb high frequencies. Similarly, reverberation that is dull sounding is associated with rooms that contain soft materials, such as plush carpets and

drapes, which readily absorb high frequencies. In this manner, reverberation imparts useful information about the composition of the surrounding space.

Reverberation is also important for establishing distance cues. In a reverberant space, when the distance between the source and the listener is increased, the level of the direct sound decreases considerably, but the level of reverberation does not decrease much. Thus, the level of direct to reverberant sound can be used as a distance cue, with dry (non-reverberant) sounds perceived as being close, and reverberant sounds perceived as being distant.

Simulating reverberation is essential for establishing the spatial context of an auditory scene. Reverberation gives information about the size and character of the surrounding space, it is very useful for correctly perceiving distances, and it adds greatly to the realism of the simulation.

3.1.10 Early Reflection Model

The early reflections in a room can be modeled by considering the walls to be acoustic mirrors. The soundfield produced by a source and a single wall reflection is equivalent to the soundfield created by the source and an "image source" located at the mirror image location behind the wall, with no wall present. Determining the early echo response for a room requires reflecting the source position across each wall boundary to obtain a set of first-order image sources, those that correspond to a single wall reflection. The free path propagation from these image sources to the listener then determines the echo response. Second-order reflections can be modeled by reflecting the first-order image sources across the wall surfaces to obtain second-order image sources. The figure below shows a rectangular room containing a source X and a listener O. Some nearby first and second-order image sources are also indicated.




Figure 3.4. Image sources in a rectangular room. The dotted line from the source to the listener represents a reflected sound path which is equivalent to the free field contribution from the indicated image source. Additional image sources are shown that correspond to other reflective paths between the source and listener. First-order image sources are labeled X1, second-order image sources are labeled X2.

Based on the positions of the source, listener, and the reflective surfaces (walls, floor, ceiling), it is easy to use the above geometrical procedure to calculate the time and direction of all early reflections. Each reflection can then be rendered using (1) a delay line to delay the sound according to the total travel time along the reflected path, (2) an attenuation or filter to approximate the transmission and reflection losses, and (3) a binaural synthesizer to properly localize the reflection.

Panorama models the first-order reflections in a rectangular room using the above procedure. There are six first-order reflections corresponding to the left, right, front and back walls, and the ceiling and floor. The ceiling and floor are called the top and bottom walls in Panorama.

Panorama models the wall reflections using filters that closely approximate the sound absorption properties of common building materials. The absorption data were obtained from standard acoustical texts.

3.1.11 Late Diffuse Reverberation

The early reflection model does not address the late portion of the reverberation, which contains millions of reflections traveling in all directions. Late reverberation is usually generated using recursive filters (filters that have feedback elements) such as comb and allpass filters.

Panorama uses a high quality diffuse reverberator based on an allpass feedback loop topology. The character of the reverberation is controlled by the following parameters: reverberation time, room size, pre-delay, and early and late damping frequencies. The reverberation time is the 60 dB decay time of the reverb. The room size parameter alters the pattern of reflections and the character of the late reverberation to simulate various room sizes. The early damping parameter controls the amount of high frequencies admitted into the reverberator. The late damping frequency parameter controls the absorption of high frequencies in the late reverberation. High damping frequencies result in a bright sounding room, low damping frequencies result in a warm sounding room.

3.1.12 Distance Cues

The principal cue for distance is the loudness of the sound. A sound source will be louder when it is closer to the listener than when it is farther away. However, this cue is often ambiguous because the listener doesn't know a priori how loud the source is. Thus, a moderately loud crashing sound could be perceived as a quiet, close crash, or a distant, loud crash. But since we know crashes are usually loud, the latter interpretation is more likely.

Another important cue for distance is the loudness of reverberation. When sound is produced in a reverberant space, the associated reverberation may often be perceived as a background ambience, separate from the foreground sound. The loudness of the reverberation relative to the loudness of the foreground sound is an important distance cue. The reason for this is due to the acoustics of reverberant spaces. The foreground sound consists largely of the sound that propagates directly from the sound source to the listener, this so-called direct sound decreases in amplitude as the distance to the listener increases. For every doubling of distance, the amplitude of the direct sound decreases by a factor of one half, or 6 dB. The amplitude of the reverberation, on the hand, does not decrease considerably with increasing distance. The ratio of the direct to reverberant amplitude is greater with nearby objects than it is with distant objects. Thus, distant objects sound more reverberant than close objects.

This relationship is diagrammed in figure 5. In the figure, the direct sound amplitude drops 6 dB for each doubling of distance. The reverberation amplitude shown in the figure starts at -12 dB relative to the direct gain and drops at 1.5 dB per doubling of distance. Hence the reverberation gain is greater than the direct amplitude at distances greater than about 6 units.

Panorama provides gain and distance rolloff parameters for editing the relationship between amplitude and distance. The gain parameter defines the starting gain in dB when the source is 1 foot away from the listener. The rolloff parameter defines the attenuation per doubling of distance. Gain and rolloff parameters are provided for the direct source, the early reflections, and the reverberation.

For the purposes of creating an effective sounding scene, it is often necessary to tweak the parameters to get the desired distance effect. For example, when synthesizing virtual acoustic scenes, it can sound unnatural if the reverberation doesn't attenuate sufficiently with increasing distance. It also becomes difficult to localize distant sound sources if there is too much reverberation. Also, attenuating the direct source by 6 dB per doubling of distance usually sounds excessive, although it is physically accurate. Natural acoustics has a very large dynamic range; when creating virtual acoustic scenes it is often necessary to compress the dynamic range by decreasing distance effects, i.e., using lower rolloff slopes than physical acoustics would dictate.

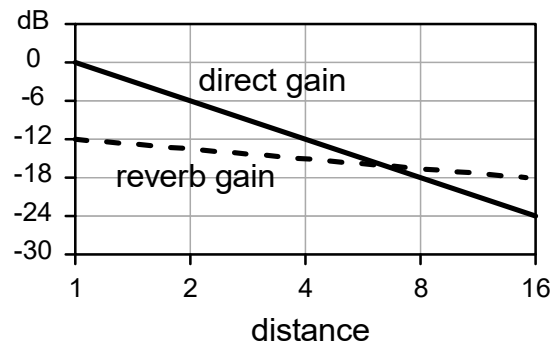


Figure 3.5. A typical distance model used by Panorama. The direct source has a gain of 0 dB and a rolloff slope of -6 dB. The reverb has a gain of -12 dB and a rolloff slope of -1.5 dB.

3.1.13 Doppler motion effect

The Doppler motion effect is commonly heard in nature as a pitch change when a speeding object passes a listener. When the object is approaching the listener, the pitch is higher than the resting pitch of the object. This is because in the time it takes the object to emit one waveform the object has moved closer to the listener, and thus the emitted wavelength is shorter than normal. Similarly, when the object is retreating from the listener, the pitch is lower than the resting pitch, because the emitted wavelengths are longer than normal.

Simulating the Doppler effect is important for generating realistic motion effects. Panorama simulates the Doppler motion effect using a variable delay line. The amount of delay is proportional to the distance between the listener and the sound object. Thus, the delay line effectively simulates the propagation of sound through the air. When the distance changes, so does the length of the delay, and the pitch also changes as it would in nature. Unlike nature, Panorama provides the feature to turn off the Doppler effect. When the Doppler effect is turned off, the amplitude changes as a function of distance, but the pitch stays constant.

3.2 Panorama Audio Routing

Following is a diagram of Panorama's audio routing:

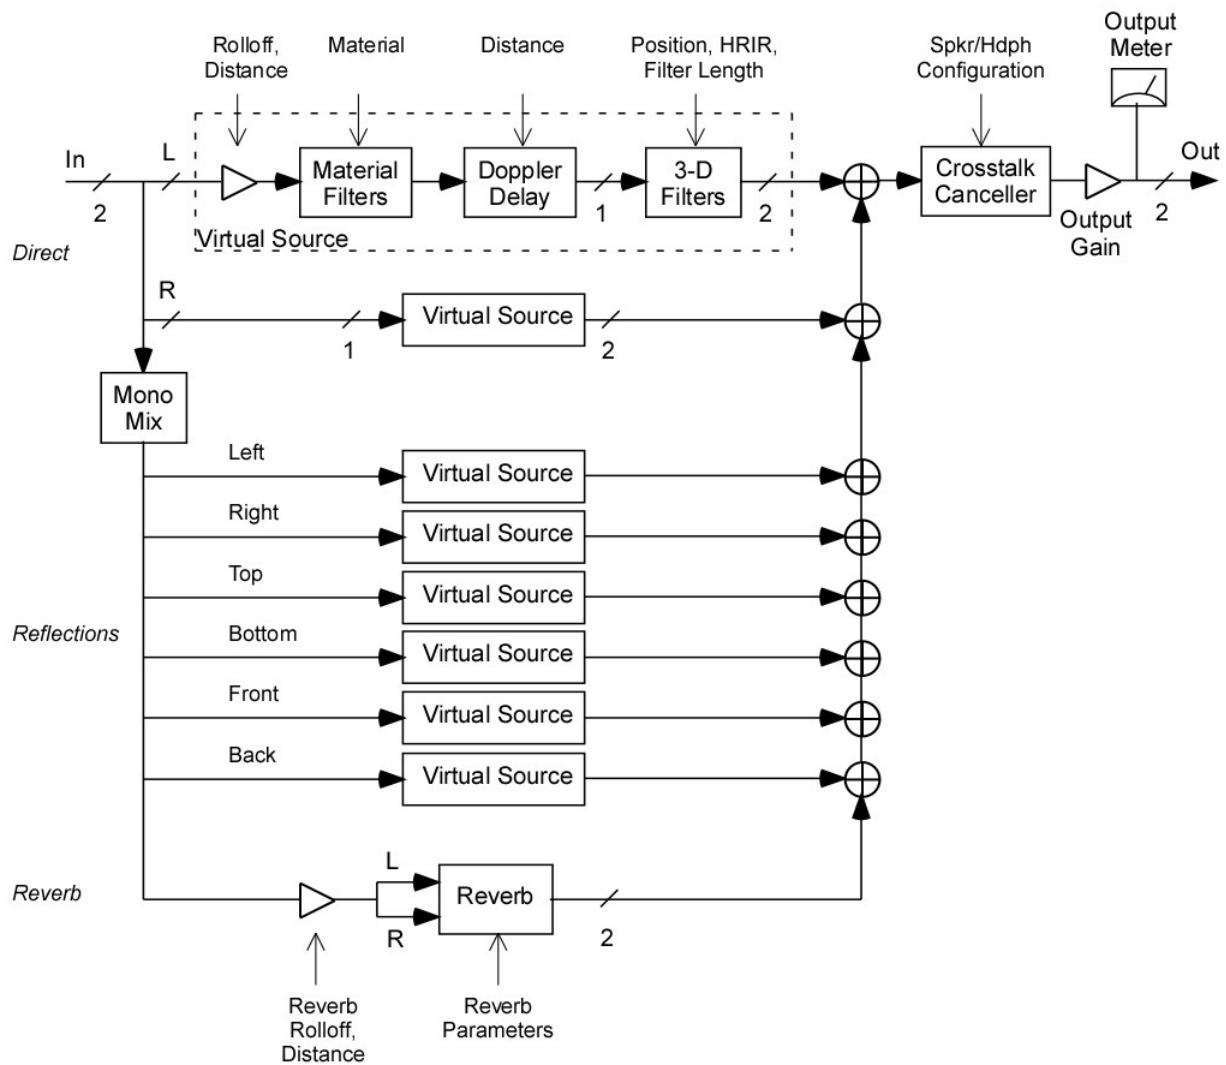


Figure 3.6. Panorama audio routing block diagram.

Panorama uses the concept of a virtual source, a monophonic sound that is reproduced in the virtual space at a particular location. The stereo input is split into left and right channels and each channel is processed as a virtual source. Virtual source processing starts with a gain stage to simulate the attenuation of sound as a function of distance from the sound to the listener. The signal is then processed by filters that simulate the reflections of sound off reflecting materials; for direct sources (those that are not reflected) this filtering is bypassed. The signal is then processed through a variable delay line that simulates the propagation of sound through air. The variable delay naturally creates the Doppler pitch effect. Finally, the signal is processed to simulate the diffraction of sound by the torso, head, and ears of the listener, labeled "3-D filters" in the figure. The resulting left and right ear signals are simulations of the signals that would reach a listeners' ears in a real listening environment. These signals are summed to a stereo output bus.

Simulating the acoustical reflections off the walls, floor, and ceiling is accomplished by creating additional virtual sources at mirror image locations behind each wall surface. The reflected virtual sources are filtered according to the sound absorption properties of the reflecting surface material. Note that the reflections are monophonic, even with stereo sources. The outputs of the reflected virtual sources are summed to the output bus.

The late, diffuse reverberation in the space is simulated using a high quality reverberator. A monophonic mix of the input is sent to reverberator and the stereo reverb output is summed to the output bus.

The stereo output bus carries the signals that should be reproduced at the ears of the listener. If loudspeaker playback is selected, the signals are processed by a crosstalk canceller. If headphone playback is selected the crosstalk canceller is bypassed. Finally, the signal is processed by the output gain stage and level metering.

3.2.1 Virtual Speakers for Stereo Inputs

Panorama uses virtual speakers to simplify the placement of stereo inputs. When processing stereo inputs, a pair of virtual speakers is created at locations to the left and right of the source position, on the same horizontal plane. The spacing of the virtual speakers is determined by the width parameter. The virtual left speaker is positioned width degrees to the left of the source and the right speaker is positioned width degrees to the right of the source. If the width parameter is set to "Mono", the stereo input is mixed to a single monophonic virtual source.

3.3 Parameters

All Panorama parameters and their value ranges are listed below, as they would be displayed by a generic parameter-value style user interface.

Parameter name	Values
X	-100 to +100 ft
Y	-100 to +100 ft
Z	-100 to +100 ft
Width	0 = Mono, 10 – 180 degrees
Direct Enable	0 = Off, 1 = On
Reflection Enable	0 = Off, 1 = On
Reverb Enable	0 = Off, 1 = On
Direct Gain	-36 to +12 dB
Reflection Gain	-36 to +12 dB
Reverb Gain	-36 to +12 dB
Direct Rolloff	-12 to 0 dB, per distance doubling
Reflection Rolloff	-12 to 0 dB, per distance doubling

Reverb Rolloff	-12 to 0 dB, per distance doubling
Direct FiltLen	1 - 256
Reflection FiltLen	1 - 256
Left Refl Enable	0 = Off, 1 = On
Right Refl Enable	0 = Off, 1 = On
Front Refl Enable	0 = Off, 1 = On
Back Refl Enable	0 = Off, 1 = On
Top Refl Enable	0 = Off, 1 = On
Bottom Refl Enable	0 = Off, 1 = On
Left Refl Dist	1 - 100 ft
Right Refl Dist	1 - 100 ft
Front Refl Dist	1 - 100 ft
Back Refl Dist	1 - 100 ft
Top Refl Dist	1 - 100 ft
Bottom Refl Dist	1 - 100 ft
Left Refl Material	0 - 17, see table
Right Refl Material	0 - 17, see table
Front Refl Material	0 - 17, see table
Back Refl Material	0 - 17, see table
Top Refl Material	0 - 17, see table
Bottom Refl Material	0 - 17, see table
Rvb Time	0.5 - 60 sec
Rvb Size	0 - 100%
Rvb EDamp	20 - 20,000 Hz
Rvb LDamp	20 - 20,000 Hz
Rvb Delay	0 - 250 milliseconds
Output Config	0 = headphones, 1 = speakers 10 degrees, 2 = 20 degrees, 3 = 30 degrees, 4 = 40 degrees
Output Gain	-18 to +18 dB
Doppler Enable	0 = Off, 1 = On
Bypass	0 = Off, 1 = On
Mode	0 = 3D panner
NHM Enable	0 = Off, 1 = On
Seq Enable	0 = Off, 1 = On
Spherical	0 = Off, 1 = On
Sync	0 = Beats, 1 = LFO
LFO Rate	0 - 200 BPM
Beat Mul	0 = 1/16..., 4 = x1..., 8 = x16
Beat Mod	0 = x1, 1 = x3/2, 2 = x2/3
Seq Dir	0 = Fwd, 1 = FwdJump, 2 = Bkwd, 3 = BkwdJump, 4 = Fwd-Bkwd
Mute Flyback	0 = Off, 1 = On
Num Slots	2 - 8
Interp Shape	0 = Linear, 1 = Exponential, 2 = Sinusoid
Shape Warp	0.0 - 1.0

Following this are eight sets of slot parameters, here are the Slot parameters for a single slot:

Parameter name	Values
SlotN Enable	0 = Off, 1 = On
SlotN X	-100 to +100 ft
SlotN Y	-100 to +100 ft
SlotN Z	-100 to +100 ft

Note that Slot1 Enable is always On.

The reflection materials are listed in the table below. The materials range from most reflective to least reflective. The "Perfect reflector" material is artificially created to have a perfectly flat frequency response.

Material Value	Material
0	Perfect reflector
1	Poured concrete
2	Plaster on brick
3	Concrete block, painted
4	Wooden floor
5	Damped plaster on brick
6	Plaster on lath
7	Glass window
8	Wooden platform
9	Concrete block, unpainted
10	Leather seating
11	Heavy carpet on concrete
12	Curtain
13	Heavy carpet on padding
14	Acoustic tile
15	Thick padding
16	Thick padding, extended
17	Plush seating
18	Audience

4. Installation and Registration

4.1 Installation

Installers for Power Suite are found on the downloads page of the Wave Arts website. There are separate installers for Mac and Windows. You can download the full suite installer or individual plug-in installers. Mac installers are “.dmg” files which after downloading will expand into a “.pkg” installer file; double-click on the “.pkg” file to launch the installer. Windows installers are “.exe” files; double-click on the “.exe” file to launch the installer. The installers provide various options for selecting which plug-in formats to install and whether to use Pace/iLok or Wave Arts licensing.

4.2 Registration

We support two licensing methods – Wave Arts licensing and PACE/iLok. When installing the plug-ins you must select which version of the plug-ins you wish to use. When you purchase a plug-in, you will be e-mailed a serial number (looks like WA-PPP-XXXX-XXXX where PPP is a product code and X is a hex digit). Use the serial number to unlock the plug-in as described below.

4.3 Wave Arts licensing

Prior to registration, the plug-ins operate in demonstration mode; they are fully functional but stop operating after 30 days. To unlock the plug-in after purchasing, open the plug-in, select the Tools->Register option, and enter your name, email address, and serial number. The plug-in will contact our registration server and download a license file which will unlock the plug-in. You should see a message saying your registration was successful.

If your computer is not connected to the internet, use the Tools->Offline Register option and follow the instructions to generate a keyfile at our website registration page and import the keyfile.

If you have purchased a plug-in suite, when you unlock any one of the plug-ins within the suite, the entire suite will be unlocked.

4.4 PACE/iLok licensing


All our plug-ins support PACE/iLok. Prior to activation, the plug-ins will allow you to start a 30-day trial by creating an iLok account. To unlock the plug-in after purchasing, go to our Product Registration page, select “PACE/iLok”,

and enter your serial number. A PACE redeem code (looks like XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XX) will be displayed and also emailed to you. There are two ways to redeem the code and generate a license. When opening the plug-in a dialog window will appear and give you the option to Activate the plug-in, you can paste the PACE redeem code there, and proceed to create or login to an iLok account and then transfer the license to an iLok or your machine. Otherwise, go to <http://www.ilok.com>, create an iLok account, and download and install the iLok License Manager. Within the manager, under the Licenses menu, select "Redeem Activation Code" and paste your redeem code. Then transfer the license to either an iLok dongle or your machine. The plug-in will run only if it can find a license on an iLok or the machine.

When purchasing a plug-in suite, the redeem code will generate multiple licenses, one per plug-in in the suite, but the licenses are grouped together.

5. Plug-in Control Operation

5.1 Knobs




Please refer to the following guide for information about the various ways you can use knobs:

Function	Mac	Windows
Increase/Decrease a parameter value (rotate clockwise/counterclockwise)	Click on the knob + drag up/down -or- Mousewheel	Click on the knob + drag up/down -or- Mousewheel
Fine adjustment — increase/decrease	Shift + click + drag up/down -or- Command + click + drag up/down	Right click + drag up/down -or- Shift + click + drag up/down -or- Control + click + drag up/down
Reset knob to default value	Double-click	Double-click

By default knobs follow up/down mouse motion, but this can be changed in the preferences settings.

5.2 Text Entry



Many value displays are editable text. A text field is editable if your mouse cursor changes to an I-beam when moved over the text. Following is a table that fully describes how to use the text editing features:

Function	Mac	Windows
Enter text entry mode	Click in the display	Click in the display
Select text	Click + drag	Click + drag
Select entire text	Double-click	Double-click

Delete character to left of cursor	Delete	Backspace
Delete character to right of cursor	Fn+Delete	Delete
Move the cursor left/right	Left/Right arrow keys	Left/Right arrow keys
Extend the current selection	Shift + click + drag -or- Shift + left/right arrow keys	Shift + click + drag -or- Shift + left/right arrow keys
Exit text entry mode	ESC -or- Return/Enter -or- Click on panel	ESC -or- Return/Enter -or- Click on panel
Select next parameter to edit	Tab	Tab
Select previous parameter to edit	Shift + Tab	Shift + Tab

You'll find that many parameters, such as frequency, will recognize units typed into the text field. The following values, when typed into a frequency value box, are equivalent:


$$2k = 2 \text{ kHz} = 2000 = 2000 \text{ Hz}$$

5.3 Sliders



Function	Mac	Windows
Increase/Decrease a parameter value	Click + drag	Click + drag
Fine adjustment — increase/decrease	Shift + click + drag	Right click + drag -or- Shift + click + drag -or- Ctrl + click + drag
Reset slider to default value	Double-click	Double-click

6. Menu Bar and Preset Manager



This section describes the operation of the menu bar, preset manager, and the other functions available in the menus.

6.1 Enable

The enable button looks like a power button. When lit the plug-in is enabled, and when unlit the plug-in is bypassed – it passes audio but without modification.

6.2 Undo/Redo

Clicking the Undo button causes the parameters and view options to revert to their settings prior to the last edit. The undo stack is unlimited, so you can keep clicking and revert back to the settings when the plug was initially opened. Clicking Redo restores the last undo. You can continue to redo back to the settings before undo was first used. However, if you make any edits the redo stack is discarded.

The A and B buffers (described below) have independent undo/redo stacks. The undo/redo stacks store all parameters and many of the UI view options. Plug-in preferences are not stored in the undo/redo stacks.

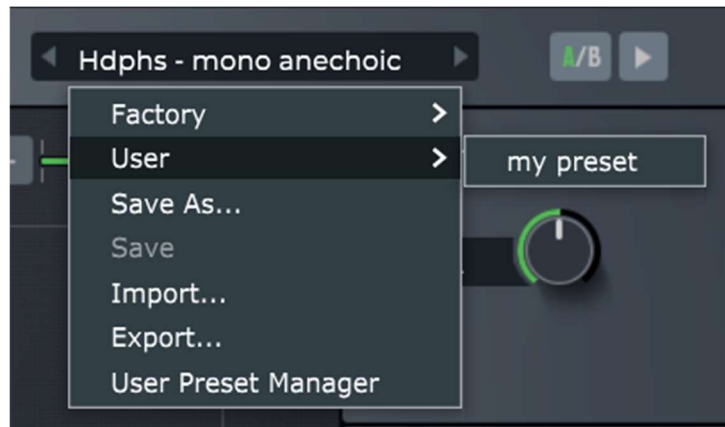
6.3 Preset name and arrow controls

The currently selected preset name is displayed in the text field in the menu bar. Changing any parameters causes an asterisk (*) to be displayed at the end of the name. This indicates that changes have been made to the preset. In order to save the changes to a user preset you must select the "Save..." item in the Preset menu, described below.

The arrow controls to the left and right of the preset name cycle through the set of factory and user presets. Clicking the right arrow goes to the next preset, clicking the left arrow goes to the previous preset.

6.4 Preset menu

The Preset menu contains lists of factory and user presets for easy selection, and options for managing presets. The functions are described in the following sections.



6.4.1 Factory Presets

Factory presets are selected from a rolloff menu at the top of the Preset menu. Factory presets cannot be modified or deleted.

6.4.2 User Presets

User presets are selected from a rolloff menu just below the Factory presets in the Preset menu. When you first run a Wave Arts plug-in, there will not be any user presets and the menu will be empty. When you save a preset using the "Save" option the preset is added to the User menu. All instances of a plug-in share the same set of user presets. So, after you save a preset with one instance of a plug-in, you can go to another instance and find that the preset can be found in its User preset menu too.

User presets are stored in a text file called "<plugin> Presets.txt", where <plugin> is the name of the plug-in you are using. If the file is deleted, an empty preset file will be created automatically the next time the plug-in runs. User presets files are stored in the following directory, depending on the operating system, where <username> is your login name:

Mac OS-X:

/home/<username>/Library/Application Support/Wave Arts/<plugin>/

Windows:

C:/Users/<username>/AppData/Roaming/Wave Arts/<plugin>/

6.4.3 Save As...

When you have created an effect you want to save as a preset, select the "Save As..." option. You will be asked to name the preset and the preset will be saved in the set of User presets. If you supply the same name as an

existing user preset, the preset will be overwritten with the new preset without any warning notice.

6.4.4 Save

The Save option is used to save changes to a user preset. The Save option is enabled when the current preset is a user preset and you have made changes, in which case an asterisk (*) will be appended to the preset name. Select the Save option to save the changes to the preset.

6.4.5 Import...

User presets can be written to files using the "Export" function, and read from files using the "Import" function. Selecting the "Import..." option will open a file chooser to select the preset file for importing. After selecting the file the presets are read and will appear in the User Presets menu. Importing does not delete existing presets; if an imported preset has the same name as an existing preset, you will get two presets with the same name!

Import can be used to import presets from an earlier version of the plug-in; the preset parameters will be migrated to the current version of the plug-in.


6.4.6 Export...

Selecting the "Export..." option will open a file chooser to specify the output file. Then your user presets are written to the file.

Preset Export is also useful for making backup copies of your user presets. If you have a large set of user presets, be sure to export them to a backup file.

6.4.7 User Preset Manager

The User Preset Manager provides a simple interface to organize user presets. You can rename presets, delete presets, and organize them by moving them up/down in the list.



6.5 A/B buffers

The A/B edit buffers allow you to compare two different sets of parameters or presets. One of the A or B indicators is always lit; the one that is lit shows the current buffer. Clicking the A/B button will switch to using the other buffer, thus changing the effect settings (assuming different settings are stored in A and B). This is quite literally an A/B compare function.

Once you have settings you like in buffer A, switch to buffer B and setup different settings, then click A/B to switch between the two.

6.6 Copy buffer

Clicking the Copy (arrow) button switches to the other buffer *while also copying the buffer*, hence the two buffers will be equal after the copy. The arrow points to the right when A is selected, and points left when B is selected.


Typical buffer use is:

- 1) Get settings you like in buffer A.

- 2) Copy to B, switching to buffer B.
- 3) Further edit the settings in buffer B.
- 4) Click A/B to compare the two buffers.

6.7 Tools menu

The Tools menu contains various important options, described below.




6.7.1 Size

The Size menu rolloff lets you change the interface size with one mouse gesture. After changing the size, clicking the plug-in title in the title bar will toggle between the changed size and 100% size.

6.7.2 About...

The About option displays important information about your plug-in. An example is shown below:



On the top line, the plug-in name and version are displayed, along with the current plug-in format (AAX, VST3, AU), bit depth, and CPU architecture. This is useful if you aren't sure which format of the plug you are running. The build date of the plug-in is displayed on the next line. If the plug-in is using Wave Arts licensing, the registration status is displayed on the next line. If the plug-in is operating in demo mode, the time remaining (if any) is displayed. If the plug-in has been successfully registered (unlocked), the serial number is displayed. If the plug-in is using Pace/iLok licensing, it will display "Pace/iLok licensing".

6.7.3 Register...

Select this option to register (unlock) your plug-in. Enter your name, email address, and serial number, and click OK.

A dark-themed dialog box titled "Register" with a red close button in the top right corner. The text inside reads: "To register, enter your serial number, name, and email address." Below this text are three input fields: "Serial #:" followed by a text box, "Name:" followed by a text box, and "Email:" followed by a text box. At the bottom of the dialog are two buttons: "Cancel" on the left and "OK" on the right.

6.7.4 Offline Register...

Select this option to register the plug-in when your computer is not connected to the internet.

6.7.5 Open User Manual...

Select this option to open this user manual in a browser.

6.7.6 Preferences...

This option opens the Preferences dialog to customize the plug-in operation.

Trademark Notices



VST is a trademark of Steinberg Media Technologies GmbH

License Agreement

END-USER LICENSE AGREEMENT FOR WAVE ARTS SOFTWARE

IMPORTANT-READ CAREFULLY: This Wave Arts End-User License Agreement ("EULA") is a legal agreement between you (either an individual or a single entity) and Wave Arts, Inc. ("Wave Arts") for the Wave Arts software accompanying this EULA, which includes computer software and may include associated media, printed materials, and "on-line" or electronic documentation ("SOFTWARE PRODUCT" or "SOFTWARE"). By exercising your rights to make and use copies of the SOFTWARE PRODUCT, you agree to be bound by the terms of this EULA. If you do not agree to the terms of this EULA, you may not use the SOFTWARE PRODUCT.

SOFTWARE PRODUCT LICENSE

The SOFTWARE PRODUCT is protected by copyright laws and international copyright treaties, as well as other intellectual property laws and treaties. The SOFTWARE PRODUCT is licensed, not sold.

GRANT OF LICENSE This EULA grants you the following rights:

Installation and Use.

When using Wave Arts licensing, you may install and use the SOFTWARE PRODUCT on up to three computers, provided you are the only user. A license must be acquired for each individual user of the SOFTWARE PRODUCT. When using Pace/iLok licensing you are granted a single user license which you can transfer to an iLok dongle or to a machine.

Reproduction and Distribution.

You may not reproduce or distribute the SOFTWARE PRODUCT except to make backup copies, or to install as provided for above.

DESCRIPTION OF OTHER RIGHTS AND LIMITATIONS

Limitations on Reverse Engineering, Decompilation, and Disassembly.

You may not reverse engineer, decompile, or disassemble the SOFTWARE PRODUCT, except and only to the extent that such activity is expressly permitted by applicable law notwithstanding this limitation.

Software Transfer.

You may permanently transfer all of your rights under this EULA, provided you retain no copies, you transfer all of the SOFTWARE PRODUCT, and the recipient agrees to the terms of this EULA.

Termination.

Without prejudice to any other rights, Wave Arts may terminate this EULA if you fail to comply with the terms and conditions of this EULA. In such event, you must destroy all copies of the SOFTWARE PRODUCT and all of its component parts.

COPYRIGHT

All title and copyrights in and to the SOFTWARE PRODUCT (including but not limited to any images, text, and "applets" incorporated into the SOFTWARE PRODUCT), the accompanying printed materials, and any copies of the SOFTWARE PRODUCT are owned by Wave Arts or its suppliers. The SOFTWARE PRODUCT is protected by copyright laws and international treaty provisions. Therefore, you must treat the SOFTWARE PRODUCT like any other copyrighted material.

LIMITED WARRANTY

NO WARRANTIES. Wave Arts expressly disclaims any warranty for the SOFTWARE PRODUCT. The SOFTWARE PRODUCT and any related documentation is provided "as is" without warranty of any kind, either express or implied, including, without limitation, the implied warranties or merchantability, fitness for a particular purpose, or noninfringement. The entire risk arising out of use or performance of the SOFTWARE PRODUCT remains with you.

NO LIABILITY FOR CONSEQUENTIAL DAMAGES. In no event shall Wave Arts or its suppliers be liable for any damages whatsoever (including, without limitation, damages for loss of business profits, business interruption, loss of business information, or any other pecuniary loss) arising out of the use of or inability to use this Wave Arts product, even if Wave Arts has been advised of the possibility of such damages. Because some states/jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

MISCELLANEOUS

If you acquired this product in the United States, this EULA is governed by the laws of the State of Massachusetts.

If this product was acquired outside the United States, then local laws may apply.

Support

For assistance, please send email to:

support@wavearts.com

If you are having problems with a plug-in, please include the following information: plug-in name, operating system, and host software you are using. The version numbers are also helpful.

For software updates, revision history, frequently asked questions (FAQ), and more, please visit our website at:

www.wavearts.com

Index

A

A/B buffers.....42
About43

B

Binaural synthesis21
Bypass.....39

C

Crosstalk.....23

D

Direction18
Distance.....28
Distance rolloff.....28
Doppler..... 11, 29

F

Factory presets40

H

HRIR 10, 20, 23
HRTF See HRIR

I

iLok34
Import preset41
Installation34

K

Knobs37

L

Licensing34

M

Materials33
Mute flyback18

O

Output.....9

P

PACE34
Presets39

R

Reflections 14, 26, 31
Register44
Registration.....34
Reverb..... 14, 25, 27
Reverberation31

S

Save preset41
Sequencer.....17
Sliders38
Slots.....17
Spatialization.....16
Spherical paths19

T

Text entry37
Tools menu43

U

Undo39
User presets40