O7A Reverb Plugins

v2.4.0

Table of Contents

1 Introduction

The O7A Reverb plugins add synthetic reverberation to seventh order ambisonic (O7A) streams. Both algorithmic and convolution reverbs are available, including a "shaped" convolver.

In general we recommend the use of O7A streams primary for research purposes, or in studios with high-end modern hardware. **The CPU load generated by the O7A plugins can be very high and may not be suitable for use on standard studio computer systems.** Third order (O3A) streams are an alternative. These carry enough spatial detail for most practical applications and present a much lighter processing load. Many O3A plugins run roughly four or sixteen times faster than their O7A equivalents.

1.1 Compatibility

O7A streams require audio busses with at least 64 channels, which (at the time of writing) most Digital Audio Workstations (DAWs) cannot handle. Some are limited at 2 channel stereo and many are limited at the 8 channels used for 7.1 surround mixes. **The O7A plugins *will not work correctly* with these DAWs**, which may even crash. In general, Reaper or Pro Tools Ultimate (version 2023.6 or later) are good options.

1.1.1 AAX

These plugins are compatible with Pro Tools Ultimate version 2023.6 or later.

Pro Tools keeps careful track of exactly what stem formats are present on tracks and busses; this is a powerful feature which helps avoid mistakes and helps ensure plugins are used in the right places.

1.1.2 VST2

The plugin library works as a "shell" plugin. This means that a number of individual audio plugin effects are provided by a single library file. Some VST2 hosts may have a slightly different way of managing these plugins to ordinary ones. For instance, in Max/MSP the vst~ plugin uses "subname" messages to specify the individual plugin within the library. At the time of writing, VST2 shell plugins are not supported in Nuendo or Cubase.

Most VST2-compatible DAWs (such as Reaper) have a plugin "path", which is a list of directories which will be searched for VST2 plugins. You may need to change this path to point at the location of the plugins, or move the plugins there. By default, these plugins are installed into `/Library/Audio/Plug-Ins/VST` on macOS. Various directories may be used on Windows, but `C:\Program Files\Steinberg\VST2` is not uncommon.

1.1.3 Buffering

Some of the O7A plugins use internal buffering with a length of 128 samples. For smooth CPU load, you may wish to ensure your DAW buffer size is a multiple of this.

1.2 Getting Started

This documentation assumes you are familiar with Reaper or a similar DAW, and with the O7A Core library, which provides a number of essential tools for working with O7A streams. If you are not, you may wish to start with the documentation for the O7A Core.

1.3 The License Manager

The Blue Ripple Sound License Manager application can be used to move license keys around between computers.
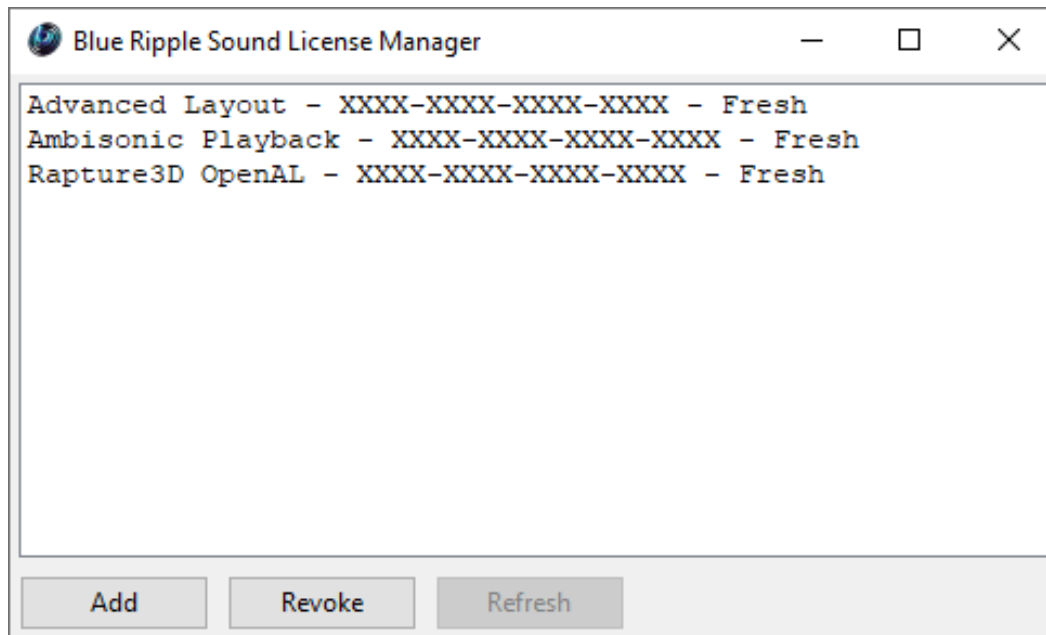
The application is generally available in the Start Menu on Windows, and in your installation directory, which is typically:

- `C:\Program Files\Blue Ripple Sound\` on Windows, and
- `/Applications/Blue Ripple Sound/` on macOS.

License keys can be removed from a computer using the 'Revoke' button and added with 'Add'. If data is not 'Fresh' this probably indicates a network connectivity issue, in which case the 'Refresh' button may be used. Generally, licenses are refreshed automatically.

1.4 Technical Requirements

1.4.1 Operating System

Supported versions of Microsoft Windows are 10 or 11, 64bit Intel.

Supported versions of Apple macOS are 10.14 to 13.5, 64bit Intel or Apple Silicon.

1.4.2 VST Host

The VST plugins require a VST 2 host with shell plugin support.

These plugins do not work with all VST 2 hosts. One reason for this is that the O7A plugins need large numbers of channels on their input or output busses. Check the individual plugins for the channel counts they need, but all of the O7A plugins need at least 64 channels on each track. Also, at the time of writing VST 2 shell plugins are not supported in Nuendo or Cubase.

1.4.3 AAX Host

The AAX plugins require **Pro Tools Ultimate version 2023.6** or later.

1.4.4 PC Hardware

Please check your PC meets the following requirements:

- Intel Core i7 CPU or better, or Apple Silicon.
- 400MB of free disk space.

1.4.5 Permissions

You'll need administrator permissions while installing on Windows. The software probably won't install properly using a "restricted" account.

1.4.6 Internet Connection Required

This software requires an Internet Connection for license activation and verification.

Successful license verification isn't required every time you use the software, but it is needed during installation and needs to succeed once every couple of weeks to keep the license fresh.

The license can be "revoked" to remove it from one machine so it can be moved on to another. You should also do this if you're updating your system in case the machine appears to have changed identity.

2 Using Reverb

2.1 Which Reverb Should I Use?

The MS3 and MS5 algorithmic reverbs both provide early and late reflections, so these are normally good places to start. They both use the same reverb algorithm internally and the main difference between them is the set of parameters available. The MS3 reverb provides just 10 parameters, which are intended for ease of use. The MS5 reverb provides around 90 parameters, which typically take longer to set up but allow more control. Both of these reverbs have a number of presets available.

An O7A Early Reflection plugin is a more elementary building block, synthesizing a single early reflection using the same algorithm as in the MS3 and MS5 reverbs.

The O7A Reverb - Convolution plugin provides simple convolution reverb, but using a 3D impulse response file, for instance recorded in a real acoustic space. The output of this reverb does not depend on the spatial directivity of its input; its directivity is as provided by the impulse response.

O7A Reverb - Shaped Convolution also uses a 3D impulse response file but gives advanced control of output directivity. This can be controlled selectively by the directivity of the input to the plugin. For instance, the reverb can be made to "bloom" from around a violinist to all around the room. In some configurations, this convolver can also produce 3D output from a mono impulse response file.

The MS3, MS5, Convolution, Shaped Convolution and O7A Early Reflection plugins all add reverb to a complete 3D O7A scene. The O7A Shoebox plugin works in a very different way. In use, it is closer to a panner, because only mono enters it but a 3D O7A soundfield leaves it. The way this happens is elaborate: a rectangular ("shoebox") room is used as an acoustic model and the plugin calculates accurate reflections of a panned sound off its walls.

This can produce highly realistic early reflections, even for moving sources, but is expensive in CPU terms and is not well-suited to generating late reflections. To add these in, we recommend sending to an MS3 or MS5 reverb with its early reflection model disabled, or a convolver with a suitable impulse response.

2.2 Effect Sends

Reverb algorithms are relatively expensive in CPU and it is normal practice to set up a single reverb to which audio is sent from a number of tracks. The amount of reverb for each track can then be controlled with a send gain. When using a reverb this way, you will probably want to set its Dry Level control to "-inf" as the dry signal will be heard anyway on the sending tracks.

For instance, an O7A track can be dedicated to an MS3 reverb. Typically, you can then send audio to it from other O7A tracks using your DAW's track routing. This approach has the additional advantage that you can apply further spatial processing to the reverb output without affecting the dry signal from the sending tracks.

When using reverb to simulate distance, it is often worth using a "pre-FX" or "pre-fade" send, as the reverb level then stays the same when the sound is faded up or down, allowing the sound to disappear into a wash of reverb.

Effect sends works very well with the scene-based reverbs such as the MS3 but is *not* a good idea when using O7A Shoebox. This acts as a panner, so everything sent to it will end up in one location! You can send the output of O7A Shoebox to another reverb of course, which can be an effective way to add a final late reverb tail.

2.3 Reverbs and Mono Inputs

The MS3, MS5, Convolution and Early Reflection plugins are all designed to be able to produce 3D reverb if fed with only plain mono inputs (rather than a full scene), if a little care is taken. This can be particularly useful when routing mono material to a reverb effect send before a complex panning FX chain. The Shoebox and Shaped Convolution plugins are not intended to be used in this way.

The mono should be sent to the first audio input of the plugin (the default in Reaper). It will then be treated an omnidirectional signal. Ensure that other inputs to the plugin are silent.

The late reflection part will then behave normally, but you may wish to pay more attention to early reflection behaviour if you are using it (this does not apply to the Convolution reverb). In particular, rotations will have no effect on the (omnidirectional) input, so to give the early reflections clear directivity you may wish to make stronger use of the 'Directional Mask' early reflection features of the O7A Early Reflection or MS5 plugins.
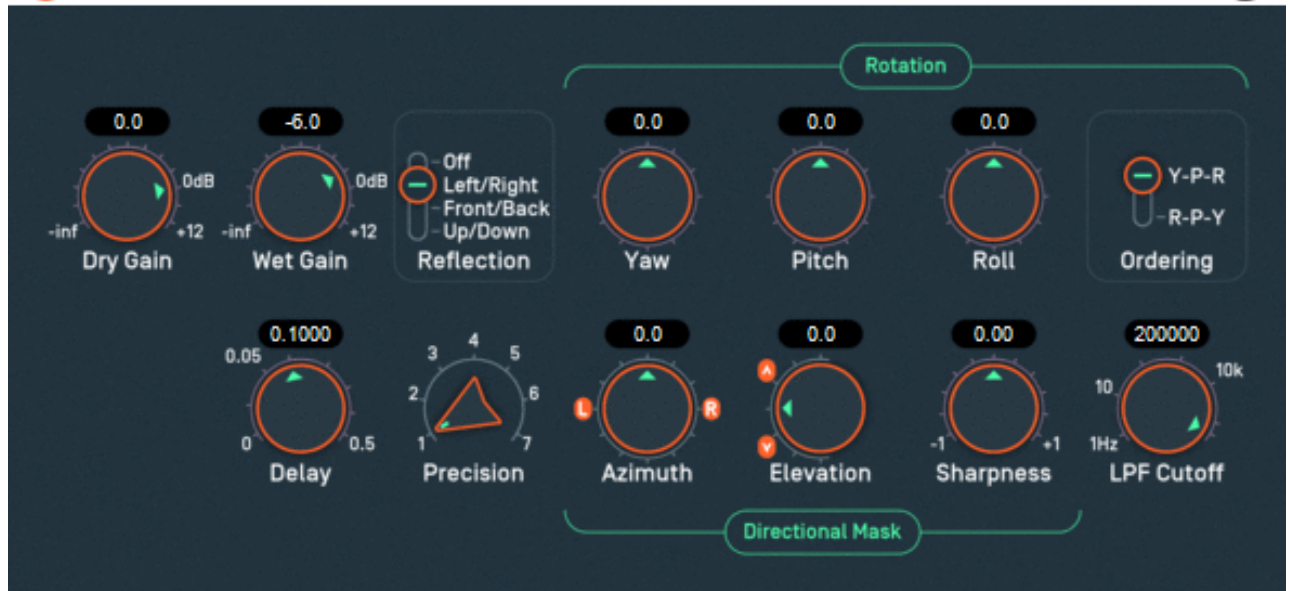
Note that when used this way, the omnidirectional input will not normally be suitable to be mixed directly into the O7A output, so the Dry Gain should be set to "-inf" (as for send plugins in general).

3 Algorithmic Reverb

3.1 O7A Early Reflection

3.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

3.1.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

3.1.3 Controls

- Dry Gain
- Wet Gain
- Reflection
- Pitch
- Roll
- Rotation Ordering
- Yaw
- Delay
- Precision
- Mask Azimuth
- Mask Elevation
- Mask Sharpness
- LPF Cutoff

3.1.4 Description

This plugin produces a single modified echo or reflection of an O7A stream.

The input may be delayed, modified and then mixed with the dry input. Modifications may include reflection and rotation, directional emphasis and low pass filtering.

A number of these reflections are used in the MS3 or MS5 reverbs, along with a late reflection model.

3.1.5 Controls

3.1.5.1 Controls: Dry Gain and Wet Gain

These control how much unmodified "dry" input and modified "wet" material is included in the output.

3.1.5.2 Control: Reflection

This can be used to apply a fixed spatial reflection to the O7A material before a spatial rotation is applied. The spatial reflections available are:

- None
- Left/Right
- Front/Back
- Up/Down

Note that, by combination with the rotation, it is possible to construct any other angle of reflection.

3.1.5.3 Controls: Pitch, Roll, Rotation Ordering and Yaw

The delayed signal can be rotated. The rotation is applied after any fixed reflection.

This modification stage behaves in the same way as the O7A Rotation plugin.

These controls use angles between -180 and +180 degrees which can each be thought of as controlling individual rotations around three fixed axes. The overall rotation is a combination of the three.

Rotation	Axis	Description of Positive Movement
Yaw	Z	The soundfield is "rotated" left around the middle of the room. For instance, a sound that starts at the front moves left, a sound that starts at the left moves backwards and a sound directly above does not move.
Pitch	Y	The soundfield is "tumbled" backwards. For instance, a sound that starts at the front moves upwards, a sound that starts at the left does not move and a sound directly above moves backwards.
Roll	X	The soundfield is "tilted" clockwise. For instance, a sound that starts at the front does not move, a sound that starts at the left moves upwards and a sound directly above moves right.

The ordering control determines the order in which the yaw, pitch and roll rotations are applied to make up a single rotation. The default is Y-P-R (i.e. yaw, pitch and then roll) but this can be switched to R-P-Y (i.e. roll, pitch and then yaw).

3.1.5.4 Control: Delay

This is a delay applied to modify the signal. It is measured in seconds.

When choosing delay times, it is sometimes useful to consider the time it takes for sound to travel a distance. The speed of sound is typically around 340m/s, so it takes about 0.003 seconds (3ms) for sound to travel one metre.

3.1.5.5 Control: Precision

Precision can be reduced to save CPU and produce a less sharp image. The precision value corresponds directly to the internal ambisonic order used for processing. A value of seven indicates that full seventh order ambisonics (O7A) will be used.

3.1.5.6 Controls: Mask Azimuth, Mask Elevation and Mask Sharpness

After reflection and rotation, the modified signal can be masked so only certain directional components are kept.

This modification stage behaves in the same way as the O7A Directional Mask plugin from the O7A Manipulator plugin library. This modifies levels in a 3D O7A mix to emphasise or de-emphasising sound in a particular direction.

Azimuth is the horizontal angle, between -180 to +180 degrees, measured anticlockwise (left) from the front. Elevation is a vertical angle between -90 and +90 degrees, measured with positive upwards and 0 on the horizontal.

For instance, the direction for azimuth +90 and elevation +45 can be found by turning 90 degrees to the left and then looking up by 45 degrees.

The sharpness control affects the sharpness and strength of the emphasis applied. A value of zero indicates that no modification should be made to the stream. Positive values leave sound in the chosen direction as it is but reduce sound elsewhere, so that the direction is emphasised. Values near +1 provide the sharpest emphasis. Negative values reduce the sound in the chosen direction, so it is de-emphasised. Values near -1 provide the sharpest de-emphasis.

3.1.5.7 Control: LPF Cutoff

This final stage of signal modification applies a first order low pass filter to the audio. This can be used to make the reflection sound more distant or muffled.

3.1.5.8 That's a lot of controls. How should I use them?

As is often the case, the O7A Visualiser will help you out. You might want to put the dry level to silent while working on the reflection.

For instance, to simulate a side reflection off a right wall, you might want to reflect the whole stage so that sounds on the left appear on the right etc. and rotate slightly for the exact wall angle. And to isolate just reflected sounds coming from the right, apply a positive directional mask on the right hand side. This will result in a reflection heard on the right which only occurs after audio events on the left.

3.2 O7A Reverb - MS3
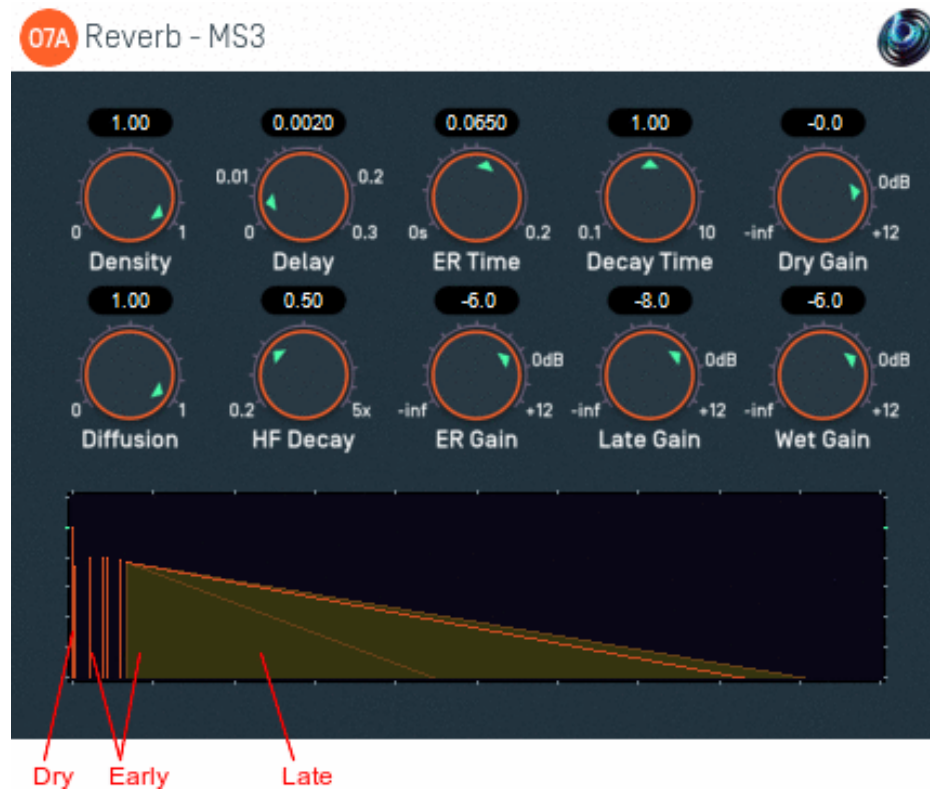
3.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

3.2.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

3.2.3 Controls

- Dry Gain
- Wet Gain
- Delay
- ER Time
- ER Gain
- Decay Time
- Late Gain
- HF Decay
- Density
- Diffusion

3.2.4 Description

This plugin adds synthetic reverberation to an O7A stream. It uses the same high quality algorithm as the MS5 reverb, but with simplified controls.

The output of the plugin includes five simulated early reflections followed by a late reflection model which provides a decaying texture of simpler reflections.

A simple visualisation of the reverb is shown, with time along the horizontal axis and decibels along the vertical. The dry (unprocessed) signal, early reflections and late reflections are shown simplistically to help set controls correctly. Once filters, diffusion, masking and so on are taken into account the resulting impulse response typically has a lower level than shown here.

A range of presets are available for this plugin.

3.2.5 Controls

3.2.5.1 Controls: Dry Gain and Wet Gain

These control how much unmodified "dry" input and modified "wet" material is included in the output. The "wet" gain applies to all reflections, both early and late.

3.2.5.2 Control: Delay

This controls the initial delay time before the reverb starts, i.e. the period of time (in seconds) between the dry output and the first early reflection.

When choosing delay times, it is sometimes useful to consider the time it takes for sound to travel a distance. The speed of sound is typically around 340m/s, so it takes about 0.003 seconds (3ms) for sound to travel one metre.

3.2.5.3 Control: ER Time

This controls the period of time (in seconds) over which the early reflections are spread.

3.2.5.4 Control: ER Gain

This controls the output level of all of the early reflections.

3.2.5.5 Control: Decay Time

This controls the overall speed of decay of the late reflection model. It is measured as the time it takes for the late reverberation to decay in level by 60 decibels. It is modified at high frequencies by the HF Decay trim control.

3.2.5.6 Control: Late Gain

This sets the output level of the late reflection model.

3.2.5.7 Control: HF Decay

This modifies the decay time of the late reflection model at high frequencies. Values below one can be used to make the late reflections sound more distant or muffled, which is generally a good idea.

3.2.5.8 Control: Density

This controls the density of the late reflections. Higher density values produce a larger number of reflections per second. Lower values can make the room sound more "ringy".

3.2.5.9 Control: Diffusion

This controls the diffuseness of the late reflections. Higher diffusion values introduce allpass filters which result in reflections with softer transients. Lower values can increase reverb "flutter".
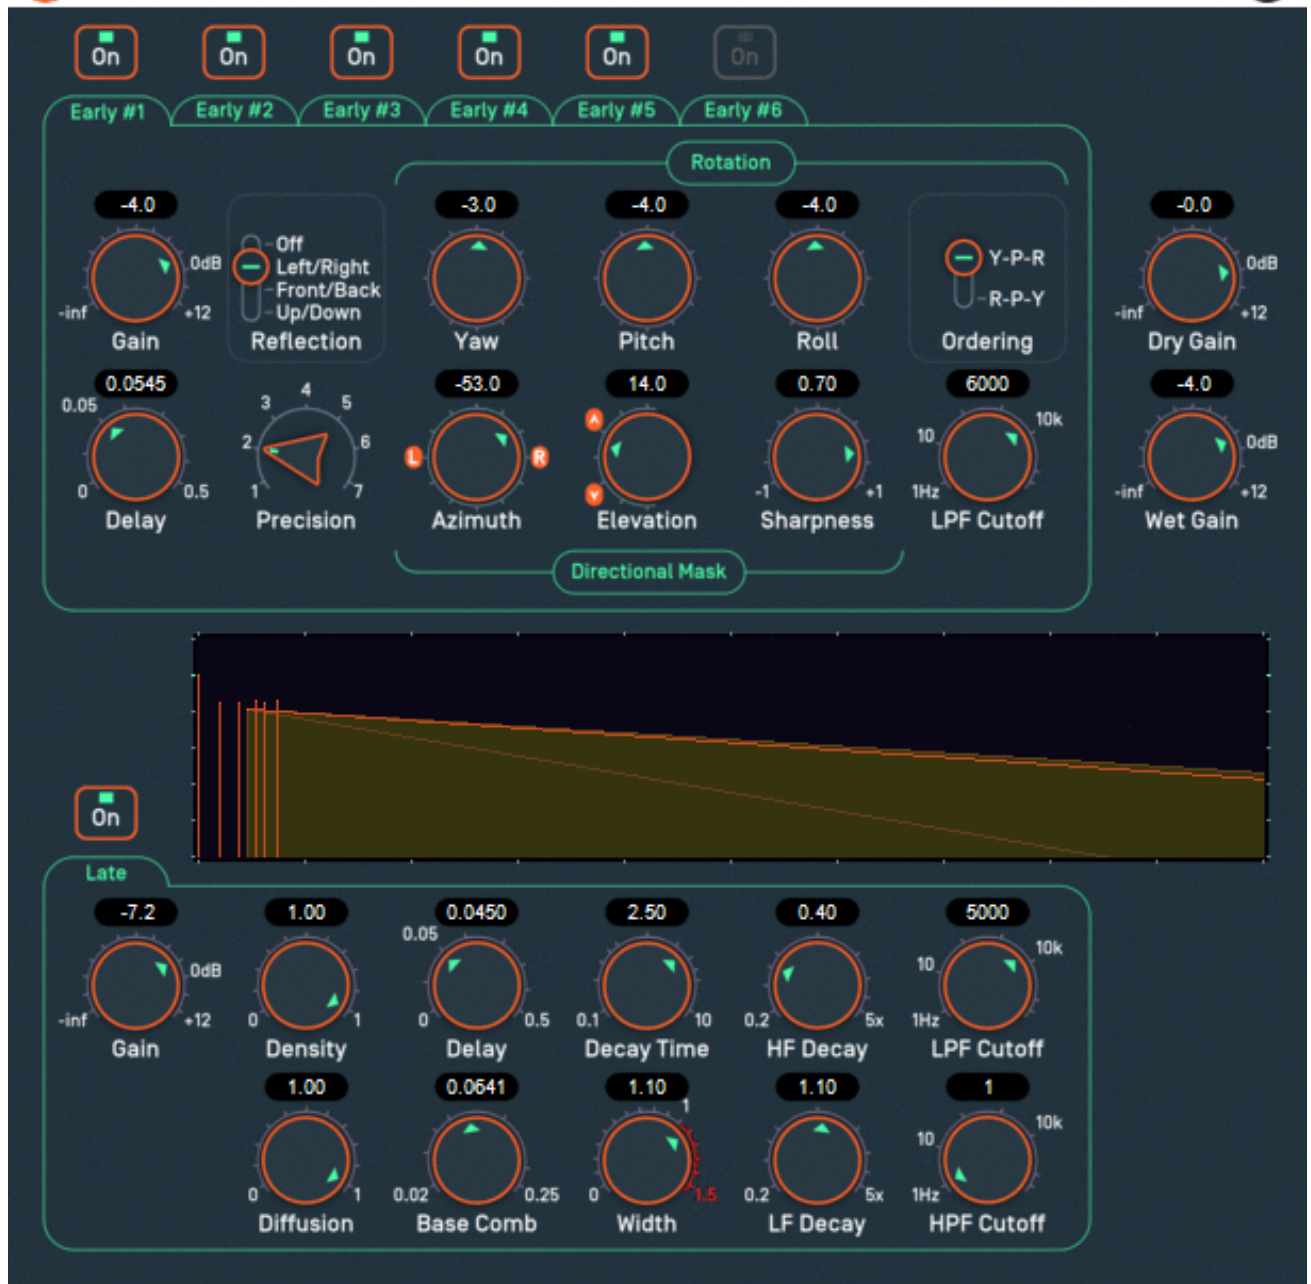
3.2.6 Presets

Presets are available for this plugin:

- Tiny
- Small
- Large
- Huge
- Scatter
- Soft

3.3 O7A Reverb - MS5

3.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

3.3.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

3.3.3 Controls

- Dry Gain
- Wet Gain
- ER1 - ER6 On/Off
- ER1 - ER6 Gain
- ER1 - ER6 Delay
- ER1 - ER6 LPF Cutoff
- ER1 - ER6 Mask Azimuth
- ER1 - ER6 Mask Elevation
- ER1 - ER6 Mask Sharpness
- ER1 - ER6 Pitch
- ER1 - ER6 Precision
- ER1 - ER6 Reflection
- ER1 - ER6 Roll
- ER1 - ER6 Rotation Ordering
- ER1 - ER6 Yaw
- Late On/Off
- Late Gain
- Late Density
- Late Diffusion
- Late Delay
- Late Base Comb
- Late Decay Time
- Late Width
- Late HF Decay
- Late LF Decay
- Late LPF Cutoff
- Late HPF Cutoff

3.3.4 Description

This plugin adds synthetic reverberation to an O7A stream. The same processing core is used as in the MS3 reverb, but with more precise control.

The plugin uses a classic early/late reflection model. Up to six precisely-controlled "early" room reflections can be simulated, along with a final tail of generic "late" reflections.

Each early reflection simulation allows the input to be delayed and modified. Modifications may include reflection and rotation, directional emphasis and low pass filtering. The controls are at the top right of the screen and are the same as in the O7A Early Reflection plugin - see that plugin for details. Each early reflection runs in parallel and each can be switched on and off independently using the six on/off switches at the top. The tabs below the switches are used to select the particular early reflection controls that are displayed.

The late reflection model provides a decaying texture of simpler reflections. This can also be switched on and off. Controls for the late reflection model are at the bottom right of the screen.

A simple visualisation of the reverb is shown, with time along the horizontal axis and decibels along the vertical. The dry (unprocessed) signal, early reflections and late reflections are shown simplistically to help set controls correctly. Once filters, diffusion, masking and so on are taken into account the resulting impulse response typically has a lower level than shown here.

This plugin does not provide directional controls for the late reflection part of the reverb. We wondered about building this into the plugin but then decided that it is quite complicated enough already. In Reaper and other software, it is possible to set up a separate reverb track and apply

spatial manipulations to the output of the reverb there.

A range of preset programs are available for this plugin. In Reaper, these appear on the main drop-down menu immediately above the main plugin screen area.

3.3.4.1 That's a huge number of controls. How should I use them?

This is a rather complicated plugin. If you want a simpler one, you might want to try the MS3 version instead, or use a preset.

One way to set up a reverb is to use a few initial early reflections to simulate distinct side, ceiling/floor and front/back reflections and then use the late reflection model to make the reverberation after them sound natural. You can mute other parts of the reverb while you are working on a specific area using the on/off controls and dry gain.

Hopefully the visualisation at the bottom of the screen will help you configure things as you would like. And, as is often the case, the O7A Visualiser will help you out with the spatial results.

3.3.5 Controls

3.3.5.1 Controls: Dry Gain and Wet Gain

These control how much unmodified "dry" input and modified "wet" material is included in the output. The "wet" gain applies to all reflections, both early and late and is cumulative with early and late reflection gains.

3.3.5.2 Controls: ER1 - ER6 On/Off

The six on/off switches at the top of the screen enable or disable the six independent early reflections individually.

3.3.5.3 Controls: ER1 - ER6 Gain

The gain control in the Early Reflection box near the top of the screen sets the output level for the selected early reflection.

If automating this plugin, there are six early gain parameters available, corresponding to each of the six early reflections, even when controls for the relevant reflection are not on screen. These are labelled "ER1Gain", "ER2Gain" etc.

3.3.5.4 Controls: ER1 - ER6 Delay, ER1 - ER6 LPF Cutoff, ER1 - ER6 Mask Azimuth, ER1 - ER6 Mask Elevation, ER1 - ER6 Mask Sharpness, ER1 - ER6 Pitch, ER1 - ER6 Precision, ER1 - ER6 Reflection, ER1 - ER6 Roll, ER1 - ER6 Rotation Ordering and ER1 - ER6 Yaw

The other controls in the Early Reflection box near the top of the screen control the currently selected early reflection. The parameters behave in the same way as in the O7A Early Reflection plugin - see that plugin for more details.

If automating this plugin, there are six versions of each early reflection parameter available, corresponding to each of the six early reflections, even when controls for the relevant reflection are not on screen.

3.3.5.5 Control: Late On/Off

This switches the entire late reflection model on or off.

3.3.5.6 Control: Late Gain

The gain control in the Late Reflection box near the bottom of the screen sets the output level of the late reflection model.

3.3.5.7 Control: Late Density

This controls the density of the late reflections. Higher density values produce a larger number of reflections per second. Lower values can make the room sound more "ringy".

3.3.5.8 Control: Late Diffusion

This controls the diffuseness of the late reflections. Higher diffusion values introduce allpass filters which result in reflections with softer transients. Lower values can increase reverb "flutter".

3.3.5.9 Control: Late Delay

This is a delay between the dry signal and the initial arrival of the late reflections. It is measured in seconds.

3.3.5.10 Control: Late Base Comb

This changes the length (in time) of repeating late reflections. Larger values imitate a larger and simpler room.

Changing this value can have timbral impact so it can be worth experimenting with. A value of 0.05 is often good.

3.3.5.11 Control: Late Decay Time

This controls the overall speed of decay of the late reflection model. It is measured as the time it takes for the late reverberation to decay in level by 60 decibels. It is modified at other frequencies by the Late HF Decay and Late LF Decay trim controls.

This control often has the greatest impact on perceived room size.

3.3.5.12 Control: Late Width

This controls the impression of spaciousness provided by the late reflections. A value of 0 reduces the late reflections to mono. Normally a value of 1 should be used here.

3.3.5.13 Control: Late HF Decay

This modifies the decay time of the late reflection model at high frequencies. Values below one can be used to make the late reflections sound more distant or muffled, which is generally a good idea.

3.3.5.14 Control: Late LF Decay

This modifies the decay time of the late reflection model at low frequencies.

3.3.5.15 Control: Late LPF Cutoff

This controls the cutoff frequency of a simple low-pass filter applied to the final output of the late reflection model. This can be used to make it sound more distant or muffled.

3.3.5.16 Control: Late HPF Cutoff

This controls the cutoff frequency of a simple high-pass filter applied to the final output of the late reflection model.

3.3.6 Presets

Presets are available for this plugin:

- Small Hall
- Medium Hall
- Large Hall
- Small Room
- Medium Room
- Large Room
- Canyon
- Cathedral
- City
- Clouds
- Cupboard
- Hills
- Standing Stones
- Stone Room
- Tunnel

3.4 O7A Shoebox
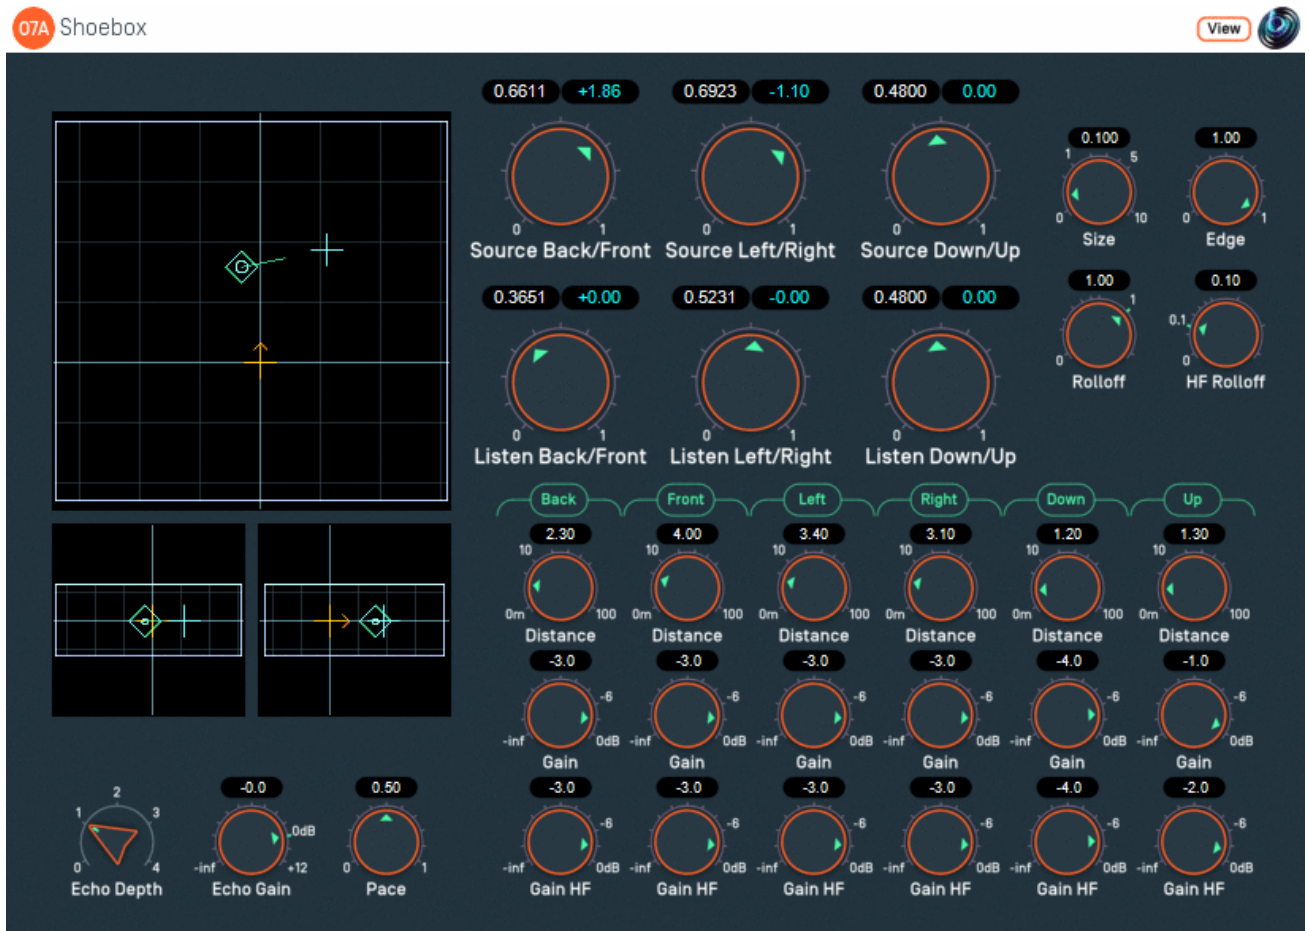
3.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

3.4.2 Audio

|	Channels	Content
Input	1	Mono
Output	64	O7A

3.4.3 Controls

- Source B/F
- Source L/R
- Source D/U
- Listener B/F
- Listener L/R
- Listener D/U
- Pace
- Source Size
- Source Edge
- Gain Rolloff
- HF Rolloff

- Echo Depth
- Echo Gain
- Wall Back
- Wall Down
- Wall Front
- Wall Left
- Wall Right
- Wall Up
- Gain Back
- Gain Down
- Gain Front
- Gain Left
- Gain Right
- Gain Up
- HF Back
- HF Down
- HF Front
- HF Left
- HF Right
- HF Up

3.4.4 Description

This plugin produces an acoustic simulation of a rectangular "shoebox" room in which a single (mono) sound "source" is playing and a front-facing "listener" is listening. The room has six surfaces, which we call walls, although the upper and lower walls would more normally be known as the ceiling and floor. The distance of the walls from a central "origin" can be changed.

The simulation calculates a number of paths from the source to the listener. The first, or "direct" path is a straight line between the source and the listener. Other paths are found by simulating reflections of the sound off the walls of the room until the listener is reached.
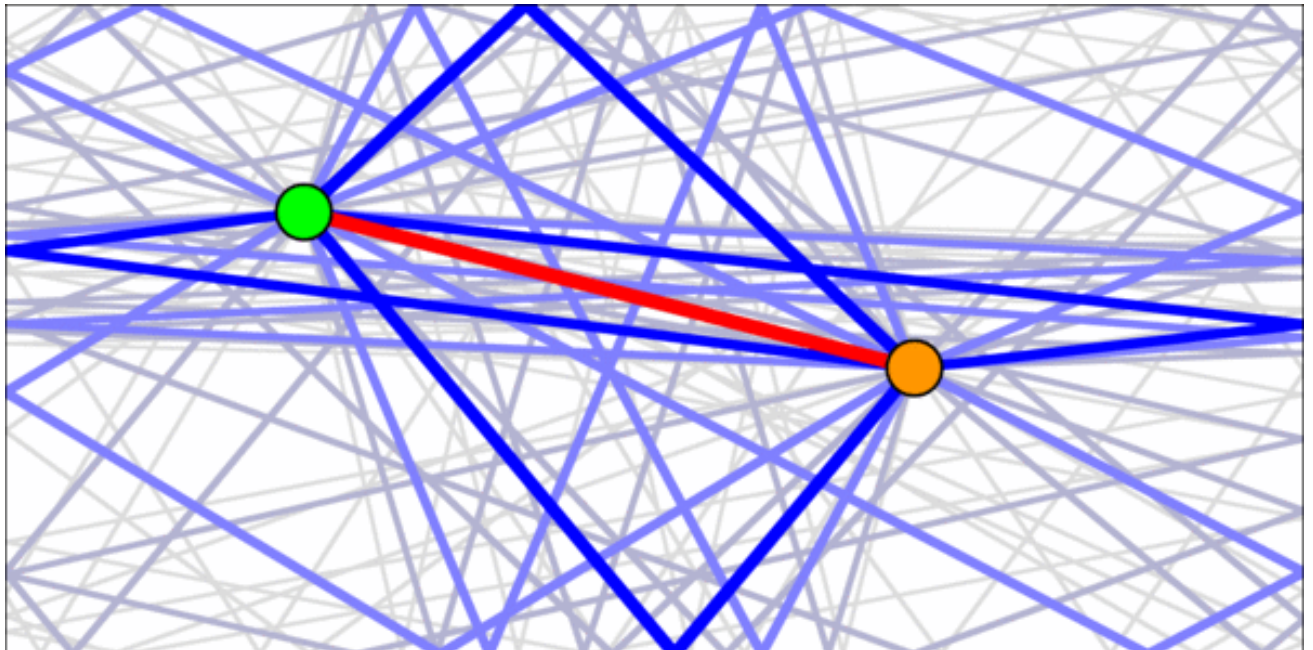
In principle, an infinite number of (increasingly quiet) reflections are needed for a strict simulation. In practice, an Echo Depth control limits the number of reflections that are simulated. You may wish to replace missing later reflections using a less precise reverb such as the MS5.

All paths (including the direct path) are modified to take into account distance attenuation (see Gain Rolloff) and high frequency air absorption (see HF Rolloff).

All paths are also delayed, taking into account the speed of sound. When the distance between the sound and listener is changing, this results in a simulation of Doppler shift, which is typically heard as a change of pitch that depends on speed.

When sounds are reflected off walls, absorption is modelled using wall-dependent gains and high frequency gains. These allow control of the "character" of the wall surfaces.

A range of preset programs are available for this plugin. In Reaper, these appear on the main drop-down menu immediately above the main plugin screen area.

3.4.4.1 The Control Surfaces

There are three control surfaces on the left side of the plugin, which can be used to move the source or listener. The larger surface is a view from above, so the front wall of the room is at the top and the left wall is on the left. The other two are views of the same scene from behind (so the ceiling is at the top and the left wall is on the left) and from the right (so the ceiling is at the top and the back wall is on the left). These views may be familiar from the O7A Panner LS and O7A Zoom plugins from the O7A Manipulators library.

The control surfaces show an outline of the room shape, with a grid. The outline changes to be proportional to the actual room dimensions.

The source is shown on the control surfaces in green/blue and the listener in orange. Left-clicking or left-dragging on the surfaces changes the "target" location of the source and the source will start moving towards this location at a speed that depends on how far away the target is and the current Pace setting. This means you can drag the source around on a control surface in a way which gives reasonable control of speed (and thus Doppler).

This listener can be moved in much the same way, but by right-clicking. If you do not have a right mouse button, you can hold down the "Alt" key on your keyboard and left-click instead.

3.4.4.2 Warning: CPU Load

This plugin can create a large CPU load, so be careful. When choosing the balance between performance and audio quality, we have chosen quality in several places where we could have used cheaper but less precise algorithms. The difference in quality is most obvious when the source or listener is moving. The plugin has been tuned to run fast, but it simply has a *lot* of processing to do.

And it gets worse. Because this plugin calculates the acoustic reflections of a source, you generally need to run a separate copy of the plugin *for each different source location*. For instance, this means that if four sounds are playing then you will typically need to run four copies of the plugin. This is different to how plugins such as the MS5 reverb work - these can be applied to the whole scene at once, rather that separately for each sound in the scene.

You may want to do a little planning before starting to use this plugin, to make sure you are going to have enough CPU power for your project.

The Echo Depth control has the largest impact on CPU load and high settings can become very expensive very quickly. Because the human brain seems to pay the most attention to the earliest reflections, a good compromise is typically to generate just the first few reflections using this plugin (using an echo depth of 1 or 2) and then to send everything to a single shared MS3, MS5 or convolution reverb to generate late reflections.

3.4.4.3 Wet/Dry Mix

In general, mixing dry signal with the output of this plugin (for instance with Reaper's Wet/Dry dial) is a Bad Idea. This is because the direct path signal produced by this plugin (which is the most similar to a 'dry' signal) is delayed, so combining it with a dry signal may result in phasing.

If you need a quick and dirty way to change the direct/reflected balance, the Echo Gain control can be used.

3.4.4.4 Listener Orientation

This plugin assumes that the listener is pointing forwards. If this is not what you want, you can modify the output of this plugin using the O7A Rotation or O7A Look plugins. You can mix the output of several shoebox models together before applying this modification.

3.4.5 View Support

When used with O7A View Sync from the O7A View library, this plugin can be connected to the separate View or ViewVR applications.

When this option is available, a 'View' button appears at the top right of the plugin's user interface.
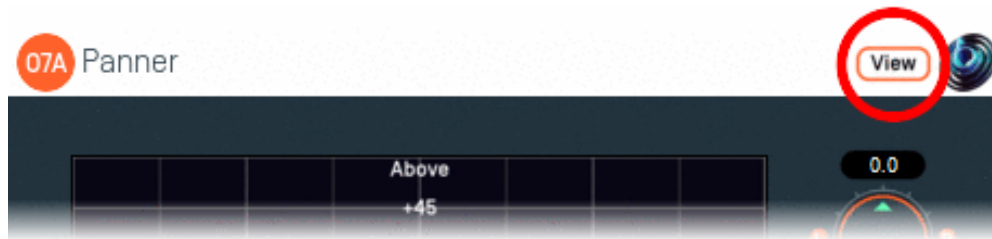
When this is pressed, the plugin is connected to the View or ViewVR application and is displayed there.

View Support	ViewVR Support
The current source direction, relative to the listener, is shown using a cross. Hold the left mouse button to set the source location.	The current source (green/blue) and listener (orange) locations are shown as locations in a scaled box, scaled using the plugin wall dimensions. Source and listener target locations are shown as 3D crosses and current locations are shown as spheres (current location movement is determined by the current Pace setting). The source location relative to the listener is also shown in its unscaled actual location. One VR controller moves the source (green/blue) and one moves the listener (orange) and they have coloured indicators to show which is which. Hold a VR controller trigger to move a scaled location to the controller.

3.4.6 Controls

3.4.6.1 Controls: Source B/F, Source D/U and Source L/R

These three controls determine where the sound source should be within the room. A scale from 0 to 1 is used for each control, which is shown in white above the dial.

Label	Meaning	Value 0	Value 1
Source B/F	Move the source to the back or front	On the back wall	On the front wall
Source L/R	Move the source to the left or right	On the left wall	On the right wall
Source D/U	Move the source down or up	On the floor	On the ceiling

An equivalent value in metres (relative to a central "origin" point) is available in blue above the dials. However, these values are not available directly through automation. Automation uses the values shown in white.

This location is also shown on the control surfaces on the left, in green/blue. It can be changed by left-clicking. Note that the source does not jump immediately to this location when it is changed. Instead, the source is pulled towards it. This allows control of the speed the sound is moving, which is important for realistic Doppler. (Also see the Pace control.)

3.4.6.2 Controls: Listener B/F, Listener D/U and Listener L/R

These three controls determine where the listener should be within the room. A scale from 0 to 1 is used for each control, which is shown in white above the dial.

Label	Meaning	Value 0	Value 1
Listener B/F	Move the listener to the back or front	On the back wall	On the front wall
Listener L/R	Move the listener to the left or right	On the left wall	On the right wall
Listener D/U	Move the listener down or up	On the floor	On the ceiling

An equivalent value in metres (relative to a central "origin" point) is available in blue above the dials. However, these values are not available directly through automation. Automation uses the values shown in white.

This location is also shown on the control surfaces on the left, in orange, along with arrow heads to show the direction in which the listener is pointing. The location can be changed by **right**-clicking. Note that the listener does not jump immediately to this location when it is changed. Instead, the listener is pulled towards it. This allows control of the speed the listener is moving, which is important for realistic Doppler. (Also see the Pace control.)

3.4.6.3 Control: Pace

When the source or listener location is set, the control surface shows the target location using a cross shape and the source or listener will start to move towards the target. On the way to the target, the current location is shown using a diamond and the direction and speed of movement is shown using a line.

Changing the pace changes how quickly the source or listener moves towards the target. By changing this and dragging the source or listener around the control surface it is possible to simulate different speeds of movement.

Values are between 0 and 1, where a value of 0 makes the source or listener move slowly and 1 quickly.

3.4.6.4 Control: Source Size

Source size can be used to give a source a volume. When the source is close, this changes the spatial response of the source (you will see it looks bigger in the O7A Visualiser). This effect continues when the listener travels inside the source.

When the listener is inside the source, there is a substantial omnidirectional component to the soundfield generated here. You may wish to make this more diffuse, for instance by using the O7A Diffuser plugin.

The source size is indicated in the control surfaces using a dark green/blue circle around the source's centre.

3.4.6.5 Control: Source Edge

The edge setting further determines the shape of a source with size. Increasing the edge setting introduces tapering from the edge of the source, effectively moving more of it towards its centre. This changes the spatial response of the source, particularly when it is close.

The source size is indicated in the control surfaces using a light green/blue circle around the source's centre.

3.4.6.6 Control: Gain Rolloff

This rolloff determines the rule that is used to reduce the source's level as the source moves away. A value of 1 is closest to how this occurs in reality. Reducing the value reduces the effect of the level drop and increasing the value exaggerates it. A value of 0 means that the level stays constant.

Mathematically, the gain applied is determined by the inverse square law, but raised to a power given by this control.

3.4.6.7 Control: HF Rolloff

This controls a simple low-pass filter that reduces high frequency content as the source moves away. This happens in reality due to air absorption. A value of 0.1 is close to reality at typical pressure, temperature and humidity, and changing this value reduces or increases the effect. A value of 0 means that no filtering occurs.

The value used with this control is the number of additional decibels lost at 10kHz for each metre of distance, on top of the main Gain Rolloff effect.

The low-pass filter used is further modified to take into account high-frequency absorption during wall reflection.

3.4.6.8 Control: Echo Depth

The echo depth is the number of times the sound is bounced off walls (between 0 and 4). Because of the number of walls involved, increasing the depth massively increases the number of audio paths considered. This allows greater realism, but also increases the CPU load, which is already high at seventh order. Load is roughly proportional to the number of paths.

Depth	Audio Paths
0	1
1	7
2	25
3	63
4	129

3.4.6.9 Control: Echo Gain

This gain (between -60dB and +12dB) affects all reflected audio paths (i.e. everything except the direct path). It can be used to make quick and dirty changes to the reverb balance, for instance to check how the direct sound on its own compares to the overall plugin output. Generally we recommend this be left at 0dB.

3.4.6.10 Controls: Wall Back, Wall Down, Wall Front, Wall Left, Wall Right and Wall Up

These values indicate how far away each of the six walls are from a central "origin" point. When these values are changed, the room layout and grids on the control surface are updated.

Distances are in metres, between 0 and 100.

3.4.6.11 Controls: Gain Back, Gain Down, Gain Front, Gain Left, Gain Right and Gain Up

These gains control how much sound is absorbed (or lost) when a sound is reflected off each wall.

Setting this to 0dB means that no sound is lost in this way. (Reflected sound will still be quieter because of the longer distance travelled.)

Setting this to "-inf" means that no sound will be reflected from the wall at all, which can be particularly relevant in outdoor simulations.

3.4.6.12 Controls: HF Back, HF Down, HF Front, HF Left, HF Right and HF Up

These gains control how much sound is absorbed (or lost) *at high frequencies* when a sound is reflected off each wall.

Setting this to 0dB means that no sound is lost in this way. (Reflected sound will still be quieter because of the longer distance travelled.)

Note that this effect is cumulative with the main gains (Gain Back etc.) above. By using one or the other, or both, different surfaces can be simulated.
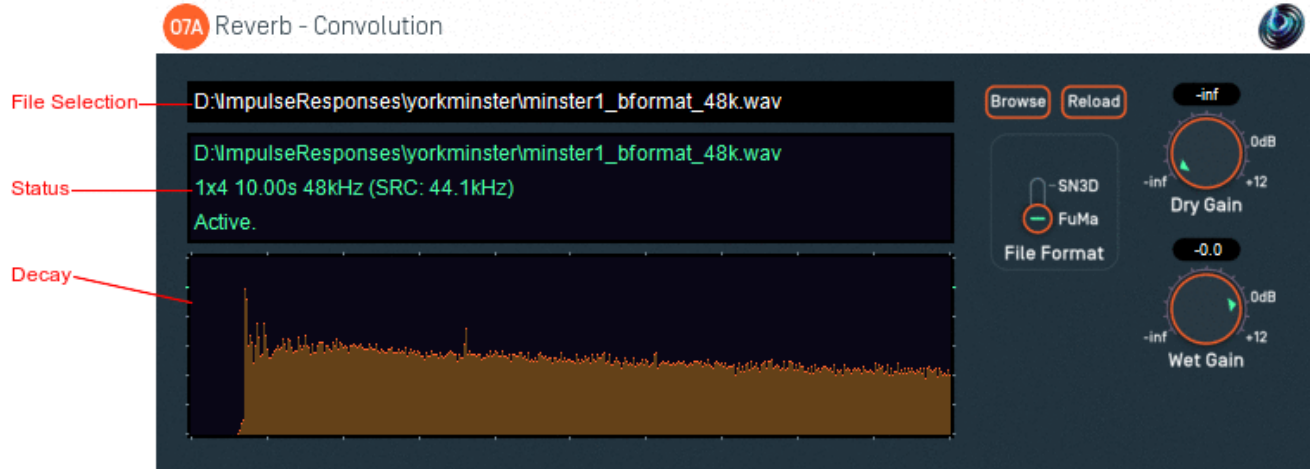
3.4.7 Presets

Presets are available for this plugin:

- Small Room
- Medium Room
- Large Room
- Soft Hall
- Hard Hall
- Park
- Valley
- Jungle
- Desert
- Cupboard
- Echoes
- Dry

4 Convolution Reverb

4.1 O7A Reverb - Convolution

4.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.1.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

4.1.3 Controls

- Dry Gain
- Wet Gain
- File Format

4.1.4 Description

This plugin provides basic convolution reverb, but designed to be easy to use with O7A streams. It is a simple plugin and does not provide impulse response editing or capture; it is assumed that you have a suitable impulse response from somewhere else (however, see the notes below on Editing Impulse Responses).

You may be familiar with convolution reverb for stereo and perhaps surround. The principles here are much the same. The essential idea is that an acoustic path in a real acoustic, or another reverb, can be captured as an "impulse response", which is the output of acoustic when it is fed with an "impulse" (i.e. a theoretical, mono, frequency-flat "click"). Capture techniques for real spaces generally use a speaker playing a sine sweep and a spatial microphone, and then calculate the impulse response from the recorded results. Cruder measurement can be performed with a starter pistol (but please get permission and warn people first!).

The plugin has a 64 channel O7A input, but this is mostly to allow some of the O7A dry signal to be mixed in to the output. The input to the convolution algorithm itself uses just the first, omnidirectional channel of the O7A and so does not depend on the spatial content of the input.

However, the output of the convolution algorithm does include spatial content, where the spatial content is derived from the impulse response. For instance, if the impulse response has a large amount of acoustic energy on the left hand side, this will be heard in the convolution results.

4.1.4.1 Impulse Responses

Impulse responses must be stored in standard PCM-encoded WAV files. The impulse response need not be at the same sample rate as the plugin; in this case, the plugin will resample the impulse response but will take significantly longer to initialise.

This plugin can use impulse responses at up to seventh order. In practice, first or second order impulse responses are often perfectly acceptable for synthetic late reverberation.

To select an impulse response, enter a path name into the File Selection box provided, or click the "Browse" button and select a file that way. If the file is changed on disk it can be loaded again with the "Reload" button.

Both SN3D and FuMa impulse responses are supported, but it is important that you tell the plugin which you are using by setting the file format.

4.1.4.2 Sourcing Impulse Responses

A number of recorded and modelled B-Format impulse responses are available from the Open Acoustic Impulse Response (Open AIR) Library at https://www.openairlib.net, under a variety of licenses.

Impulse responses can also be generated by acoustic modelling software. Realistic impulses responses can be generated for real or imaginary buildings. Software that can do this includes:

- CATT-Acoustic - http://www.catt.se/
- Odeon - https://odeon.dk/

4.1.4.3 Memory Use and CPU Load

Convolution reverb is typically much more expensive in terms of memory use and CPU load than "algorithmic" reverbs (such as the MS3). Further, the CPU load is typically roughly proportional to the length of the impulse response used, so long impulses responses can be very costly indeed. Be careful what you use as an impulse response!

You may wish to consider setting up convolution reverbs on separate tracks and sending to them, so that multiple sounds can be processed by the same reverb.

4.1.4.4 Editing Impulse Responses

This plugin does not include a sound editor to edit the impulse response. Instead, we recommend you edit the impulse response externally as an ordinary O7A sound file and press the "Reload" button to bring those changes into the plugin.

For instance, you can reduce the length of the impulse response by editing the overall envelope (shape). With a little careful mixing it is often possible to increase the length too. Applying equalization to the impulse response will change the frequency response of the reverb.

The impulse response itself can be thought of as an O7A file and can be created and/or modified using a variety of O7A plugins. For instance, you might use O7A upmixers to convert a 5.1 impulse response to O7A, apply spatial equalization, or use manipulations to change the acoustic perspective.

Impulse responses often include the "direct path" from the source to the receiver (microphone). This can be useful when convolution reverb is used to simulate sound from one specific location, but it tends to be less useful when a general reverberation is wanted. In these cases, you may wish to edit out the original direct path to avoid conflicts with other direct (dry) signals, for instance from panners or the O7A Shoebox plugin. In the latter case, or if you are also using the O7A Early Reflection plugin, you may also wish to remove or fade some of the early reflections in the impulse response to ensure they do not conflict.

4.1.4.5 Hybrid Reverb

You may also wish to consider using a "hybrid" approach to reverb, for instance using the O7A Shoebox plugin for detailed dynamic early reflections for individual sources, before sending to a shared convolution reverb for a realistic "late tail". If you do this, you may wish to remove or fade early parts of the impulse response.

Another option is to use the O7A Reverb - Shaped Convolution plugin below.

4.1.4.6 Input Signal and Levels

The "dry" signal is passed through directly, at the given input level, for instance when reverb is being added directly to an O7A mix.

The "wet" (convolved) part of the output of this plugin is not derived in a way that depends on the input directivity. Only the mono, or omnidirectional, component (the first channel) of the O7A is used and the others are ignored.

As only the first (omnidirectional) channel is used to derive the wet signal, this plugin can be used to produce reverb from a mono signal if that is placed in the first channel of the O7A input. Basic O7A panners will place audio into this first channel at the correct level. Used in this way, the dry level should normally be set to "-inf".

4.1.5 Controls

4.1.5.1 Controls: Dry Gain and Wet Gain

These control how much unmodified "dry" input and modified "wet" material is included in the output.

4.1.5.2 Control: File Format

The file format should be set to match the style of ambisonics used in the impulse response file. Setting this incorrectly may give strange results. The options are:

- "SN3D" (in the ACN channel ordering) is the format used by the O1A, O3A and O7A plugins. It is used by some recent first order microphones and is sometimes known as "AmbiX".
- First order "FuMa" is the same as classic WXYZ B-Format, which dates back to the 1970's. It was used in versions of these plugins prior to version 2.0 and some hardware and software uses it, including a number of 3D microphones. FuMa is only defined up to third order.

4.2 O7A Reverb - Shaped Convolution

4.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.2.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

4.2.3 Controls

- Dry Gain
- Wet Gain
- Env Delay
- Env Time
- Env Start
- Env End
- Locality
- Precision
- File Format

4.2.4 Description

This plugin has the features of the simpler O7A Reverb - Convolution plugin, but adds more. If you are not familiar with that plugin, we suggest you read its documentation before proceeding here, because things here are not simple.

This reverb has additional pre-processing controls that allow the directivity of its output to be "shaped" by the directivity of the live input signal. This is an extremely powerful feature, but quite mind-bending. You may wish to read this section quite carefully - but we think you will find it

worthwhile!

This plugin can create extreme CPU loads. Please use it with care.

4.2.4.1 The Directivity Shaping Envelope

The difference between this and the O7A Reverb - Convolution plugin is the presence of a "directivity shaping envelope" which runs *along the length of the reverb impulse response.* In other ways, this is a conventional envelope with a delay, a start value, a transition time and an end value. It is shown in the user interface, overlaid on the impulse response. Directivity shaping values are between zero and one.
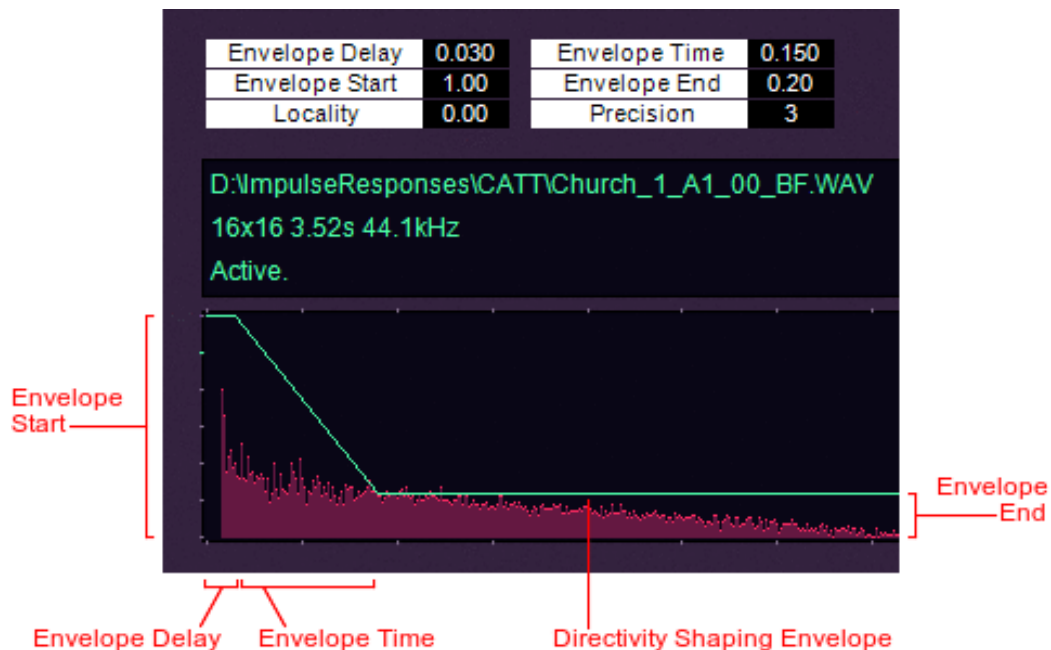
When the directivity shaping value is zero, the spatial output of the reverb is the same as with the simpler O7A Reverb - Convolution reverb, so the directivity of the output is entirely determined by the content of the impulse response file.

When the directivity shaping value is one, the output still has the timbral character and reverb decay of the impulse response file, but its directivity is overridden by the directivity *of the live input signal* at the time it was fed into the reverb.

The envelope allows the directivity shaping to change as the reverb tail progresses.

4.2.4.2 Why on Earth would I want that?

Without directivity shaping, conventional ambisonic convolution reverb will produce a wet reverb with directivity determined entirely by the impulse response.

For the later stages of the reverb, this can produce a very natural and realistic reproduction of a real space, with detailed reverberation coming from appropriate directions.

However, at the start of the reverb, the impulse response will typically contain highly directional elements relating to the original direct sound path and early reflections. These again can be highly realistic, but the directivity is fixed by the impulse response file. For instance, if the impulse response file includes a recorded direct path from the left, the reverb will always reproduce a direct path from the left, regardless of the directivity of the live input. This can be limiting when the reverb is to be added to a mix which already contains detailed sound placement.

Directivity shaping changes this. It allows the impulse file directivity to be overridden by the directivity of the live input. Further, this override can be controlled at different times along the impulse response, so it changes as the reverb decays.

Directivity shaping is typically applied strongly at the start of the impulse response, fading away to zero over a small fraction of a second. This allows the first, direct path part of the impulse response to be overridden by the "correct" directions present in the live input, but leaves later stages of the reverb decay unmodified. Slower transitions can result in a "blooming" effect starting from the original live input directions but ending with the full ambient directivity of the impulse response.
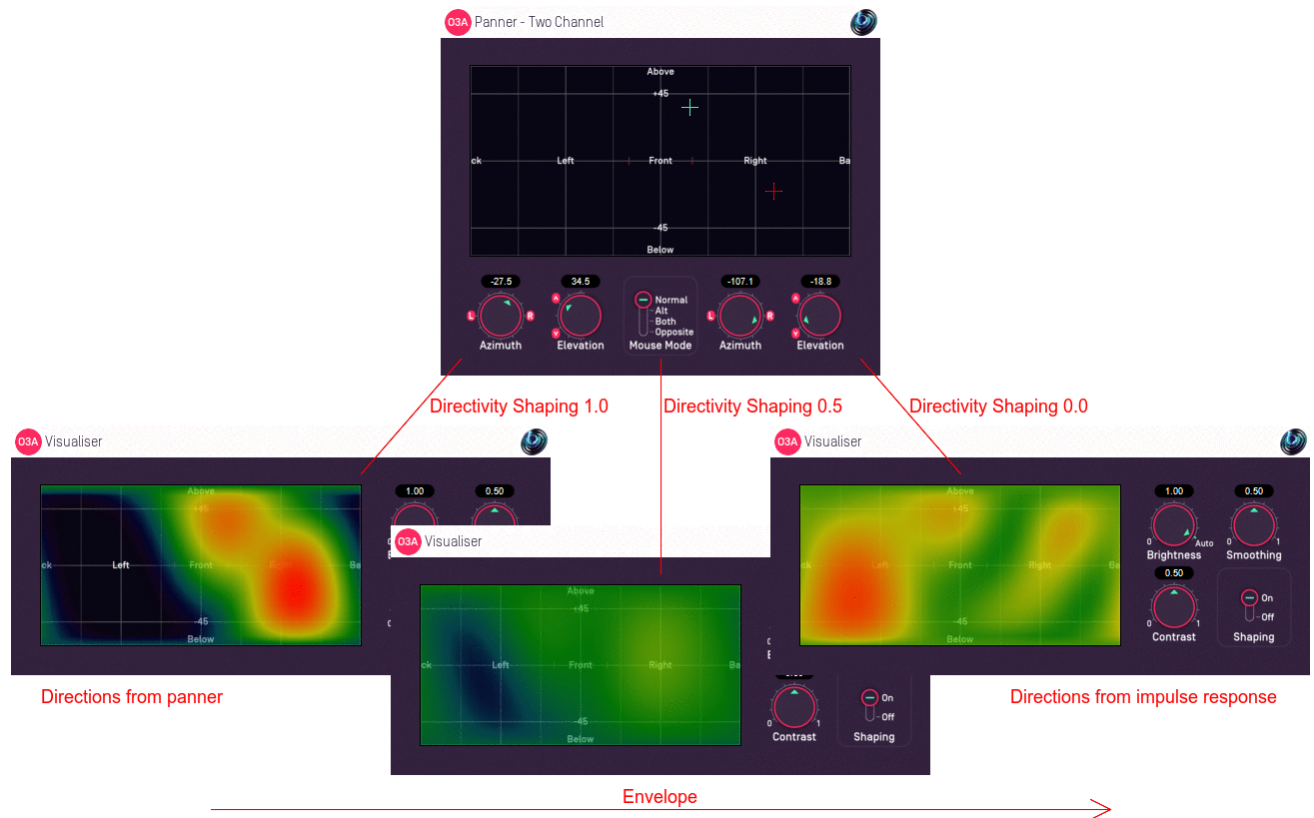
If early reflections are already present in the live input, for instance due to use of the O7A Shoebox or O7A Early Reflection plugins, these early reflections will also retain their position when directivity shaping is in use. This can be an effective way to combine dynamic early reflections with recorded late reverb.

The directivity shaping value does not have to end at zero. For instance, if it is set to one for the entire duration of the impulse response, the resulting reverberation will decay in directions driven entirely by the live input signal. This can be produce quite interesting results and, because the directivity is determined entirely by the live input, will work with just a mono impulse response file.

A number of presets (without impulse response filenames) are available to help configure the plugin.

The envelope controls are not intended as real-time automation parameters. Changing them will result in the impulse response file being loaded and processed again, which can take a while.

4.2.4.3 Memory Use and CPU Load

This plugin can use even more memory and CPU resources than O7A Reverb - Convolution, so be careful how you configure it.

Directivity shaping results in much higher resource use than with the simpler convolver. To keep this under control, it is worth looking at the directivity shaping envelope and ensuring that the envelope value is exactly zero as much as possible. This can generally be achieved by **setting the "Envelope End" control to zero** and ensuring that the "Envelope Delay" and "Envelope Time" are not set too high.

Setting "Locality" to exactly zero or exactly one will also save resources in some situations.

The "Precision" control can be used to save resources by limiting the spatial detail of the reverb output; but listen carefully to ensure the loss of spatial detail is acceptable. In particular, running at seventh order can introduce extreme CPU load and may not run in real time.

Also note that pre-processing the impulse response file takes much longer if its sample rate does not match that of the DAW project you are working with. This will show in the Status panel as an "SRC" item. This can make a big difference to usability if you are changing the pre-processing parameters repeatedly for a long impulse response.

4.2.5 Controls

4.2.5.1 Controls: Dry Gain and Wet Gain

These control how much unmodified "dry" input and modified "wet" material is included in the output.

4.2.5.2 Control: Env Delay

This is the number of seconds that the directivity shaping control stays at its start value before starting to change. This time is measured along the reverb impulse response.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.3 Control: Env Time

This is the number of seconds that the directivity shaping control takes to transition from its start value to its end value. This happens after the delay period above. This time is measured along the reverb impulse response.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.4 Control: Env Start

This is the directivity shaping control value during the envelope delay period. During the transition period the value fades from this start value to the end value.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.5 Control: Env End

This is the directivity shaping control value after the envelope transition period. During the transition period the value fades from the start value to this end value.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.6 Control: Locality

When directivity shaping is in use, the reverb direction is affected by the input signal. In these cases, the character of the reverb (timbre, decay and so on) can be set to be the same for all input directions or to vary depending on how the reverb varied spatially in the impulse response file. Locality values are between zero and one, with a value of zero giving no variation and a value of one giving maximum variation.

For example, assume that directivity shaping is in full use so the output directivity is determined entirely by the live input. Then, assume the reverb tail in the impulse response file is shorter and brighter on the left than the right. In this case, when the locality is set to one, the reverb will be shorter and brighter for sounds on the left in the input signal. When the locality is set to zero, the reverb length and timbre will be the same for all directions.

This can be even more mind-bending than other aspects of this plugin and does not make a noticeable difference with some impulse response files. If in doubt, we recommend setting the locality control to exactly zero as this can significantly reduce memory and CPU use.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.7 Control: Precision

Precision can be reduced to save CPU and produce a less sharp image. The precision value corresponds directly to the internal ambisonic order used for processing. A value of seven indicates that full seventh order ambisonics (O7A) will be used.

This control triggers pre-processing and should not be changed in real-time.

4.2.5.8 Control: File Format

The file format should be set to match the style of ambisonics used in the impulse response file. Setting this incorrectly may give strange results. The options are:

- "SN3D" (in the ACN channel ordering) is the format used by the O1A, O3A and O7A plugins. It is used by some recent first order microphones and is sometimes known as "AmbiX".
- First order "FuMa" is the same as classic WXYZ B-Format, which dates back to the 1970's. It was used in versions of these plugins prior to version 2.0 and some hardware and software uses it, including a number of 3D microphones. FuMa is only defined up to third order.

4.2.6 Presets

Presets are available for this plugin:

- Mono->Ambi, Ambi IR
- Quick Bloom, Ambi IR
- Slow Bloom, Long Ambi IR
- Directional, Mono IR
- Directional, Ambi IR
- Gentle, Ambi IR
- Tilt, Ambi IR

5 Appendix: O7A Streams

5.1 What is an O7A Stream?

A seventh order ambisonic (O7A) stream is made up of 64 individual channels of audio which together represent a 3D soundfield. Into this "audio scene" can be placed individual sound sources, reverberation and complex spatial textures.

These 64 channels can be quite confusing to understand conceptually. For instance, they do not relate to particular speaker directions, or to individual sound sources in the soundfield. You do *not* need to understand them to use them! But, it doesn't hurt to know the basics.

Each channel adds spatial detail to a sound scene. With just the first channel, you have a basically omnidirectional (mono) sound image. The second channel adds some basic detail left/right, the third up/down and the fourth front/back; these four channels make up first order ambisonics. The first and second channels together provide essentially the spatial detail available with the Mid/Side (M/S) stereo recording technique, which captures a sound image with left/right width. But with the further two channels the detail is available in all directions, not just left/right.

That covers the first 4 channels. The other 60 add further detail to make the image sharper. If you are interested in what exactly is in these channels (which is not so easy to describe) you may want to read up on Higher Order Ambisonics (HOA) and the mathematics of the Spherical Wave Equation and Spherical Harmonics. But you should *not* need to read up to *use* the techniques described here. If you want to make sense of what is going on spatially in an O7A stream, we find that it is normally best to use an O7A Visualiser or O7A Flare plugin.

5.2 What processing can I apply to an O7A stream?

As well as processing designed specifically for O7A streams, it is possible to mix streams together in the expected way. It is also possible to run conventional mono DSP algorithms on them directly by applying the algorithm to all 64 channels individually, subject to some rules. Failing to follow these rules is likely to shred the spatial imaging, so be careful! Specifically:

- If you process an O7A stream with a conventional mono DSP algorithm you *must apply the same processing to all 64 channels identically*.
- Only *linear* processing can be used (e.g. not distortion or compression). Also, be aware that time-variant processing can cause issues.

If you are working in Reaper, current versions do not make it particularly easy to set up this sort of processing, but it is possible if you are prepared to use some advanced features and do some rather tedious set-up. For instance, you can save an "FX Chain" in which a number of equalization plugins are "routed" correctly and in which controls are linked by "parameter modulation" so that one set of controls operates the others. Once the FX Chain is saved, you can load it into other projects.

5.3 Encoding

There are a number of ways in which the channels can be defined in HOA. To a large extent it does not matter which is chosen, as long as *everything uses the same convention*. This is critical and horrible things will happen to the spatial image if this is not the case. However, if you follow the convention, or convert explicitly where you need to, you will be able to pass audio around between different software packages.

These plugins use the "SN3D" ambisonic convention. As is usually the case, we order the channels using "ACN" ordering.

This encoding is used in the "AmbiX" file format and YouTube. It is supported directly by Rapture3D Universal.

5.4 How does O7A SN3D relate to FuMa and Classic Ambisonics?

The ambisonic format used (SN3D) is not directly compatible with classic 1970s four channel "WXYZ" B-Format, or the extension (FuMa) which was used by versions of these plugins prior to version 2.0.

However, conversion is straightforward. The O7A Core pack includes plugins to do this: O7A Decoder - FuMa and O7A Injector - FuMa.