

O3A Reverb Plugins
v2.4.0

Copyright 2023 Blue Ripple Sound Limited

Table of Contents

1 Introduction.....	1
1.1 Compatibility.....	1
2 Getting Started.....	2
2.1 Working with O3A.....	2
2.2 Getting Started with Reaper.....	3
2.3 Getting Started with Premiere Pro.....	8
2.4 Getting Started with Pro Tools.....	10
2.5 Getting Started with Pyramix.....	13
2.6 Audio for Virtual Reality 360 Video.....	15
2.7 Spotting A Common Issue.....	17
2.8 The License Manager.....	18
2.9 Technical Requirements.....	19
3 Using Reverb.....	20
3.1 Which Reverb Should I Use?.....	20
3.2 Effect Sends.....	20
3.3 Reverbs and Mono Inputs.....	21
4 O3A Reverb - Algorithmic Reverb.....	22
4.1 O3A Early Reflection.....	22
4.2 O3A Reverb - MS3.....	25
4.3 O3A Reverb - MS5.....	28
4.4 O3A Shoebox.....	33
5 O3A Reverb - Convolution Reverb.....	40
5.1 O3A Reverb - Convolution.....	40
5.2 O3A Reverb - Shaped Convolution.....	44
6 O3A Core - Panners.....	49
6.1 O3A Panner.....	49
6.2 O3A Panner - Classic.....	51
6.3 O3A Panner - Hemisphere.....	53
6.4 O3A Panner - Large.....	55
6.5 O3A Panner - Two Channel.....	57
6.6 O3A Panner - Eight Channel.....	60
7 O3A Core - Visualisers.....	64
7.1 O3A Visualiser.....	64
7.2 O3A Visualiser - Hemisphere.....	66
7.3 O3A Colourizer.....	68
7.4 O3A Flare.....	70
8 O3A Core - Rotation.....	72
8.1 O3A Look.....	72
8.2 O3A Rotation.....	74
9 O3A Core - Decoders and Converters.....	77
9.1 O3A Decoder - Stereo.....	77
9.2 O3A Decoder - Mono.....	79
9.3 O3A Decoder - 5.1 Basic.....	80
9.4 O3A Decoder - FuMa.....	81
9.5 O3A Injector - FuMa.....	82
9.6 O3A Decoder - O1A and O2A.....	83

Table of Contents

9 O3A Core - Decoders and Converters	
9.7 O3A Injector - O1A and O2A.....	85
10 O3A Core - Virtual Microphones.....	87
10.1 O3A Beamer.....	87
10.2 O3A Virtual Microphone.....	89
11 O3A Core - Metering.....	91
11.1 O3A Meter - Signal.....	91
11.2 O3A Meter.....	92
12 O3A Core - Soundfield Modification.....	94
12.1 O3A Gain.....	94
12.2 O3A Spatial Delay.....	95
13 Appendix: O3A Streams.....	99
13.1 What is an O3A Stream?.....	99
13.2 What processing can I apply to an O3A stream?.....	99
13.3 Encoding.....	99
13.4 How does O3A SN3D relate to FuMa and Classic Ambisonics?.....	100
13.5 Why Third Order?.....	100

1 Introduction

The O3A Reverb plugins add synthetic reverberation to [third order ambisonic \(O3A\) streams](#) by processing an existing stream, or simulating early reflections for a single source. Both algorithmic and convolution reverbs are available, including a ground-breaking "shaped" convolver.

This documentation also covers the O3A Core.

1.1 Compatibility

O3A streams require audio busses with at least 16 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 O3A plugins *will not work correctly with these DAWs***, which may even crash.

1.1.1 AAX

Most of the O3A plugins exist as AAX versions which are compatible with [Pro Tools Ultimate](#) or Studio and plugin support (or not) is indicated in this text.

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. However, where formats are not available in Pro Tools, corresponding plugins are generally not available. To avoid this in some cases, some plugins "misuse" stem formats with matching channel counts. Cases like this are described for individual plugins below in the text.

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 O3A 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.

2 Getting Started

2.1 Working with O3A

When working with stereo, you normally "pan" a mono sound to stereo and "mix" it onto a stereo bus. With O3A, you can work in a similar way but there is an extra intermediate multichannel audio format, and a final "decoding" step. The simplest case is probably as follows:

1. Panning - mono sources are panned to 16 channel O3A. The [O3A Panner](#) plugin can be used for this.
2. Mixing - multiple streams of 16 channel O3A can be mixed together using your DAW's mixing engine. No special plugin is needed here, as long as the DAW can handle 16 channel busses. You end up with a final mix, as an O3A stream.
3. Decoding - during actual playback (and you probably want to listen while you work!) you "decode" the O3A stream for whatever speaker layout you are actually listening with. One of the huge advantages of ambisonic techniques is that the O3A mix is independent of its final playback environment, so you can switch between stereo speakers, headphones, 5.1, 22.2 or whatever. One of the simplest of these decoders is provided by the basic [O3A Decoder - Stereo](#) plugin.

These are just the basics. There are other ways to encode audio into O3A rather than panning, for instance by upmixing, or recording with ambisonic microphones.

You can also manipulate O3A streams. As a simple example, it is possible to [rotate](#) the soundfield almost perfectly, and much more complex transformations are possible, including spatial ones. Manipulator plugins typically have 16 channel O3A as both an input and output. There are some [technical notes on O3A streams](#) at the end of this document.

If you are also working with first order B-Format using the SN3D/ACN convention you can mix this in too, because the first four channels of O3A correspond to the four channels of first order. To reduce an O3A stream to first order, you can simply take these four channels and throw away the rest.

2.2 Getting Started with Reaper

This section gives an introduction to use of the VST2 version of the O3A plugins with [Reaper](#).

We are assuming that you know how to use Reaper already. A number of simple Reaper example projects are included in standard installations, along with simple explanatory text files.


2.2.1 Shell Plugins

The O3A VST plugin libraries are VST2 "shell" plugins which each provide a number of plugin effects. If you find your list of O3A effects is short, with names like "O3ACore", this means that Reaper has not scanned the content of the libraries.


To fix this, go into the Reaper "Options" menu and select "Preferences..." to bring up the preferences screen. Then, under "Plug-ins", find "VST". On the VST tab, please ensure that "Get VST names/types when scanning" is checked and then click "Clear cache/re-scan".

2.2.2 Basics

In Reaper, whenever you create a track or send to be used with [O3A streams](#), make sure it has at least 16 channels. A few plugins need even more (e.g. a decoder for [Hamasaki 22.2](#)).



2.2.3 Layout



On the folder track, add the following plugins *in this order* to the FX chain:

- [O3A Visualiser](#)
- [O3A Decoder - Stereo](#) (or whichever decoder best suits your speaker rig).


This means that after all the audio from the individual parts has been mixed together it will be visualised in 3D and then converted to stereo for monitoring.

On tracks for mono parts that are to be panned, add the following:

- An [O3A Panner](#) plugin.

Do *NOT* use the built-in Reaper stereo panner with O3A!

It is also possible to add azimuth and elevation controls to Reaper on these channels. To do this, add the panner as above and then use the "Trim" button and cross the "UI" boxes for the azimuth and elevation controls.



Parts that are already first order ambisonic B-Format can typically be injected directly without issue, although processing can help (e.g. the O3A First Order Injector plugin from the O3A Upmixers set, or the O3A Harpex Upsampler). Of course, other O3A material can be injected directly. FuMa material needs to be converted using the [O3A Injector - FuMa](#) plugin.

2.2.4 Processing O3A streams

Other plugins (for instance the [rotation](#) plugin) manipulate the O3A stream. The stream needs to be in O3A format (for instance by panning) before such processing can be used.

Applying processing to the O3A stream can be confusing at first. We suggest you spend a few minutes getting familiar with the [O3A Visualiser](#) plugin and put visualisers before and after your processing. You can then see what is going on, which can be really helpful!



Some plugins have O3A as an input but output something else. In these cases, make sure you only chain together processors in a combination that make sense. For instance, the [O3A Beamer](#) plugin produces mono which you might want to run through a mono distortion plugin before panning it back into O3A using the [O3A Panner](#).

2.2.5 Sends

Effect sends can be used in the normal way. Just be careful about which busses are expecting mono or O3A as input and output and be sure to check your channel counts.

Again, do not use the Reaper panner for the send (leave it set to the centre) or the spatial image will be damaged. For similar reasons, make sure the project's "Pan Law" (in Project Settings) is set to 0dB, or avoid use of Post-Fade sends.

2.2.6 Mastering

2.2.6.1 Keep Your O3A Master Mix

Keep your Reaper project files and inputs if you like, keep your favourite stereo and 5.1 downmixes etc. But whatever you do, *keep a copy of the 16 channel O3A mix*. This may not make a difference to you now, but when someone asks you for a 22.2 (or other) downmix in a few years time, you will not need to fake height information from the 5.1 mix, or go searching for a missing file needed by the original Reaper project. Instead, you will be able to produce the downmix from the ambisonic B-Format (i.e. O3A) master using whatever tools are in vogue then.

If you have taken the simple approach as described above, probably the easiest way to grab the O3A mix is simply to switch off the decoder on your Reaper master track, and then use the "Render..." menu item. You will probably need to set the Channels to "16" by typing the number in

and use the "Stems (selected tracks)" option rather than what Reaper considers the master mix, in which case you will also need to have selected your folder track. If you want to check the output, you can always load it into a new Reaper session and run it through your favourite decoder.

2.2.6.2 Other Mixes

If you want to produce a number of different format outputs from your O3A master (stereo, headphones, 5.1 etc.) you could use a number of projects that use the O3A master file.

Another option is to place a number of decoders on your decoding/folder track and only switch on the one you want for the particular rendering.

However, our preferred option is to send from the O3A master bus to one of a number of separate decoding tracks which can be freely muted or unmuted. This is the approach taken in the example Reaper projects provided.

2.2.7 Troubleshooting

If you find plugin libraries are not picked up during Reaper plugin scanning, this can typically be fixed by backing up any relevant files and settings, removing the relevant libraries from the Reaper path (either by moving the files or changing the path), restarting Reaper, putting them back, and restarting Reaper again.

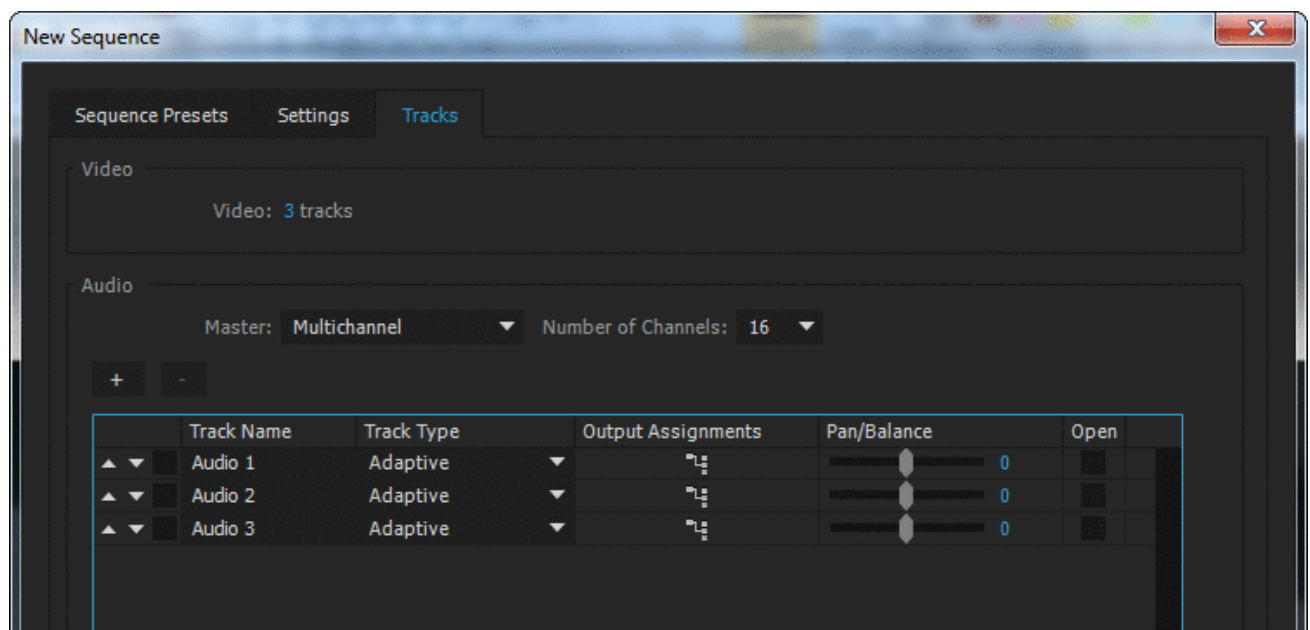
2.3 Getting Started with Premiere Pro

In Adobe Premiere Pro, as with Reaper, it is essential to make sure that your audio tracks have 16 channels (or more if you are using plugins that need them).

2.3.1 Sequence Settings

In particular, when creating a New Sequence, go to the "Tracks" tab and make the following changes:

1. The "Master" Track should be set to "Multichannel".
2. The "Number of Channels" should be set to "16", or as many as you need.
3. Each audio "Track Type" that is going to carry O3A audio (or suchlike) should be set to "Adaptive" or "Adaptive Submix".



Premiere Pro includes some VR Sequence Presets that support ambisonics. However, these only set the number of channels to 4 by default, which is correct for first order ambisonics, but not for third order. 16 channels are needed by these plugins, so please check the "Number of Channels" as above, or plugins may not work and spatialisation quality may be seriously compromised.

2.3.2 Importing Audio

When importing audio with many channels, Premiere Pro may place each channel of the audio onto a different track rather than keeping the spatial mix together. If this happens, find the clip, right click and select "Modify" and then "Audio Channels..." and do the following before trying again:

1. Set the "Clip Channel Format" to "Adaptive".
2. Check the "Active Channels Per Clip" matches the audio (for instance, 4 for first order ambisonics or 16 for third order ambisonics).
3. Set the "Number of Audio Clips" to 1.

You can also make this the default behaviour by going into Premiere Pro's "Preferences...", finding the "Audio" section and setting the "Multichannel Mono Mode" item to "Adaptive".

2.3.3 Compatibility

Our plugins do not work with all versions of Premiere. They do not work with versions prior to v12.1, or versions from v14.4 to the current version.

The "O3A View Sync" plugin from the O3A View library is not fully compatible with the current version of Premiere Pro. Necessary VST2 timing information is currently not available in Premiere, which means that video playback cannot synchronize correctly.

The Premiere list of "Available Plug-Ins" only shows one plugin from each of the plugin libraries, but they all should be available for use.

Plugin presets are not available in Premiere Pro.

2.4 Getting Started with Pro Tools

This section gives an introduction to use of the AAX version of the O3A plugins with [Pro Tools Ultimate](#), which was previously known as Pro Tools HD. Ambisonic support has been available in Pro Tools HD since version 12.8.2. It is also available in Pro Tools Studio, but not Pro Tools Intro or Artist.

These notes assume you are familiar with Pro Tools already. A number of simple Pro Tools example projects are included in standard installations, along with simple explanatory text files.

2.4.1 Stem Formats

When effects are used on tracks, Pro Tools takes account of changes in the stem format in use, which is generally very convenient when working with ambisonics. In particular, effect lists shown will be appropriate to the current stem format you are using.

For instance, if you have a mono track that you wish to pan to O3A, you will find the insert effect list includes the [O3A Panner](#) (under "Sound Field"). Introducing this effect as an insert will change the track stem format from this point to third order ambisonics (O3A) automatically. Mono effects can still be used *before* the panner, but after panning the effect list will include a different set of plugins appropriate for a third order ambisonic stream, such as the [O3A Visualiser](#) and O3A decoders. If you decode to stereo, the stem will change to stereo and so on.

2.4.2 Stem Formats that Pro Tools does not support directly

However, Pro Tools does not have direct support for some of the more unusual stem formats used by some plugins in the Blue Ripple range. Consequently:

- Some of the plugins do *not support* AAX. For instance, plugins requiring the Hamasaki 22.2 stem format are not included in the AAX version of the plugin libraries.
- Some plugins "misuse" other stem formats with matching channel counts. For instance, the [O3A Decoder - FuMa](#) plugin outputs ambisonics using the FuMa convention, but labels its output stem format as ordinary ambisonics using the SN3D convention, which is not really correct. As a result, the Pro Tools effect lists shown will not be correct. Similarly, the [O3A Decoder - IMAX 6.0](#) plugin outputs audio targeting the IMAX 6.0 layout, but labels its output stem format as standard 6.0, which does *not* have exactly the same channel definitions.


Where you are using standard Pro Tools stem formats only, these issues are not relevant. But, if you are using stem formats that Pro Tools does not support directly, care is needed.

Please see the documentation for each individual plugin for details.

2.4.3 Ambisonic Tracks

Pro Tools stem formats make it very straightforward to organise an ambisonic project. Often, tracks do not need to be created in special way as simply including an O3A panner or other ambisonic encoder will result in the track stem format changing automatically, e.g. from mono.

However, if you wish to create an ambisonic track directly, Pro Tools provides simple options on the track creation screen.



Like these plugins, Pro Tools normally uses the [SN3D/ACN](#) ambisonic format. You can convert from FuMa to SN3D using the [O3A Injector - FuMa](#) plugin.

2.4.4 Organising Projects

A recommended way to organise a project is to route all of your ambisonic material (panned or otherwise) to a single ambisonic "O3A Master" bus, set to third order ambisonics in much the same way as tracks (see above).




This gives a convenient place to bring the whole O3A mix together and send it to a number of tracks for visualisation (e.g. with [O3A Visualiser](#) or [O3A Flare](#)), metering and decoding for specific output formats.

For instance, you might create a "Stereo Decoding" track including the [O3A Decoder - Stereo](#) plugin, so the final stem format is stereo, ready for monitoring. You might also set up a "5.1 Decoding" track in parallel and mute the one you are not using - and so on for other final output formats.

Effect sends can be used in the normal way and Pro Tools will help you match up the stem formats in use.


2.5 Getting Started with Pyramix

This section gives an introduction to use of the VST2 version of the O3A plugins with [Pyramix](#). We are assuming you know your way around Pyramix already and have configured it to find the VST plugins. Note that modern Pyramix requires the 64bit version of the plugins.



2.5.1 Setup

[O3A streams](#) are most easily handled using Multi Channel Strips (MCS). When creating these, set the channel count to 16 (or even more, e.g. for [Hamasaki 22.2](#)).



Custom 16 channel mix busses can also be used.

Create General Mixing Bus

Number of Busses:

Bus Type: Mix Bus Aux Bus Object Bus
 Mix Group Aux Group

Channels Mapping:

Number of Channels:

Room Type: Virtual
 Sized [m]

Create Cancel

2.5.2 Mixing

Set mix busses to "Router" rather than "Pan" to ensure that the O3A plugins are determining the panning used.

In Pyramix, plugin user interface dials are set to the VST "linear" style, so drag up and down rather than around to change them.

Beyond this, plugins can be used much as you would expect!

2.6 Audio for Virtual Reality 360 Video

The [O3A streams](#) used by these plugins capture a spatial audio scene that can be used for Virtual Reality (VR) playback. Typically, the audio is delivered to the VR device in a raw ambisonic format and is rotated using head tracker data before live decoding. This has the effect of modifying the audio scene to take into account where the user is looking. Because raw ambisonic audio is delivered to the device, final mixes should not generally be decoded in the studio, except for monitoring.

O3A streams contain third order SN3D/ACN (AmbiX), which uses 16 channels. This, or a subset of the channels, is directly compatible with many VR systems.

2.6.1 YouTube 360-Degree Videos

YouTube 360-degree videos can use first order SN3D. To reduce your mix to first order, just take the first four channels of your 16 channel mix, or use [O3A Decoder - O1A](#).


2.6.2 Google Jump Inspector

Google Jump Inspector can use the third order SN3D audio produced by these plugins directly.


2.6.3 Facebook 360 Video

Facebook 360 videos can use second order SN3D audio, which uses nine channels.

The easiest way to reduce an O3A mix to just nine channels in Reaper is often to set the render channel count to nine by *typing* the number 9 into the Channels box.



Once you have a nine channel second order file, this can be encoded using the Facebook FB360 Encoder (v1.3.0 or later). Select format "B-format ambiX 2nd order".



2.6.4 Rapture3D Universal

The Blue Ripple Sound Rapture3D Universal game/VR engine can render the third order SN3D audio used by these plugins directly.

2.6.5 Quad Binaural

Quad binaural is an alternative system that does *not* use raw ambisonic audio. Instead, four binaural stereo images are used, corresponding to different head orientations (front, left, back and right). These are mixed together live depending on where the user is looking. This system only supports one axis of rotation, unlike ambisonics which supports all three.

Quad binaural is generally delivered as four stereo files. To produce these, we recommend use of the the O3A Decoder - Quad Binaural plugin from the O3A Decoding pack. This will typically give better final results than rotating the O3A stream and using a more conventional binaural decoder to generate the Quad Binaural.

2.7 Spotting A Common Issue

In many DAWs it is easy to forget to set your track or send channel counts to 16 at some point and end up trying to handle O3A as stereo audio. We certainly do this regularly!

This is such a common problem that it is worth taking a moment to look at what then happens, so you can spot the issue more quickly:

- Sound positioning becomes vague.
- Sound movement is only to the left and right, with nothing definite in other directions.
- The O3A Visualiser typically shows a vague left or right image, or just one colour. The left-hand case typically looks something like this:




The reason this happens is that O3A needs 16 channels to work properly. If only the first two channels are present, these happen to be the omnidirectional and left-right ambisonic components, like in M/S stereo. These are enough to get an approximate sound level to transmit some basic left/right information, but definitely not the full detail of the soundfield!

So, keep your eyes and ears open and remember that O3A needs 16 channels.

2.8 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.

2.9 Technical Requirements

2.9.1 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.

2.9.2 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 O3A 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 O3A plugins need at least 16 channels on each track, and some need more. Also, at the time of writing VST 2 shell plugins are not supported in Nuendo or Cubase.

2.9.3 AAX Host

The AAX plugins require Pro Tools Ultimate or Studio.

2.9.4 PC Hardware

Please check your PC meets the following requirements:

- Intel Core i5 CPU or better, or Apple Silicon.
- 200MB of free disk space.

2.9.5 Permissions

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

2.9.6 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.

3 Using Reverb

3.1 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 [O3A 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 [O3A 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.

[O3A 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 O3A Early Reflection plugins all add reverb to a complete 3D O3A scene. The [O3A Shoebox](#) plugin works in a very different way. In use, it is closer to a panner, because only mono enters it but a 3D O3A 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.

3.2 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 O3A track can be dedicated to an MS3 reverb. Typically, you can then send audio to it from other O3A 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 [O3A Shoebox](#). This acts as a panner, so everything sent to it will end up in one location! You can send the output of O3A Shoebox to another reverb of course, which can be an effective way to add a final late reverb tail.

3.3 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 as an omnidirectional signal. Ensure that other inputs to the plugin are silent.

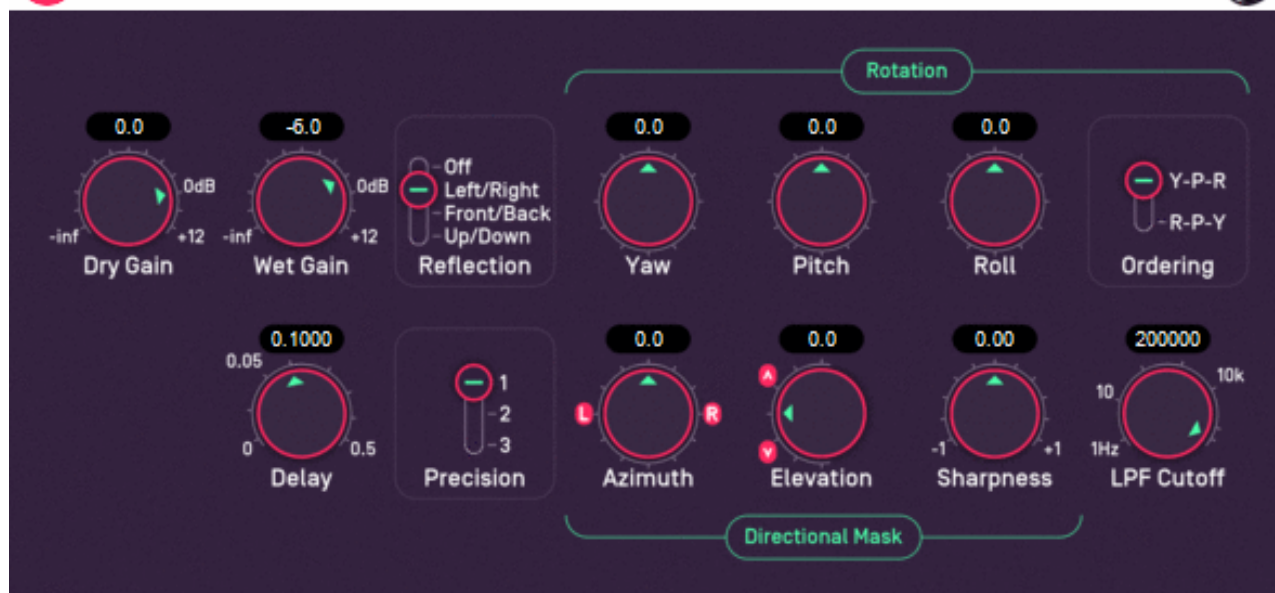
The late reflection part will then behave normally, but you may wish to pay more attention to early reflection behaviour if you are using it (this does not apply to the Convolution reverb). In particular, rotations will have no effect on the (omnidirectional) input, so to give the early reflections clear directivity you may wish to make stronger use of the 'Directional Mask' early reflection features of the [O3A 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 O3A output, so the Dry Gain should be set to "-inf" (as for send plugins in general).

4 O3A Reverb - Algorithmic Reverb

4.1 O3A Early Reflection

O3A Early Reflection



4.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

4.1.3 Controls

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

4.1.4 Description

This plugin produces a single modified echo or reflection of an O3A 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.

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: Reflection

This can be used to apply a fixed spatial reflection to the O3A 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.

4.1.5.3 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 [O3A 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).

4.1.5.4 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.

4.1.5.5 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 three indicates that full third order ambisonics (O3A) will be used.

4.1.5.6 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 O3A Directional Mask plugin from the O3A Manipulator plugin library. This modifies levels in a 3D O3A 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.

4.1.5.7 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.

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

As is often the case, the [O3A 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.

4.2 O3A Reverb - MS3



4.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

4.2.3 Controls

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

4.2.4 Description

This plugin adds synthetic reverberation to an O3A 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.

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. The "wet" gain applies to all reflections, both early and late.

4.2.5.2 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.

4.2.5.3 Control: ER Time

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

4.2.5.4 Control: ER Gain

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

4.2.5.5 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.

4.2.5.6 Control: Late Gain

This sets the output level of the late reflection model.

4.2.5.7 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.

4.2.5.8 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".

4.2.5.9 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".

4.2.6 Presets

Presets are available for this plugin:

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

4.3 O3A Reverb - MS5

O3A Reverb - MS5



4.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.3.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

4.3.3 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

4.3.4 Description

This plugin adds synthetic reverberation to an O3A 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 [O3A 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.

4.3.4.1 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 [O3A Visualiser](#) will help you out with the spatial results.

4.3.5 Controls

4.3.5.1 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.

4.3.5.2 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.

4.3.5.3 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.

4.3.5.4 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 [O3A 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.

4.3.5.5 Control: Late On/Off

This switches the entire late reflection model on or off.

4.3.5.6 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.

4.3.5.7 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".

4.3.5.8 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".

4.3.5.9 Control: Late Delay

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

4.3.5.10 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.

4.3.5.11 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.

4.3.5.12 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.

4.3.5.13 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.

4.3.5.14 Control: Late LF Decay

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

4.3.5.15 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.

4.3.5.16 Control: Late HPF Cutoff

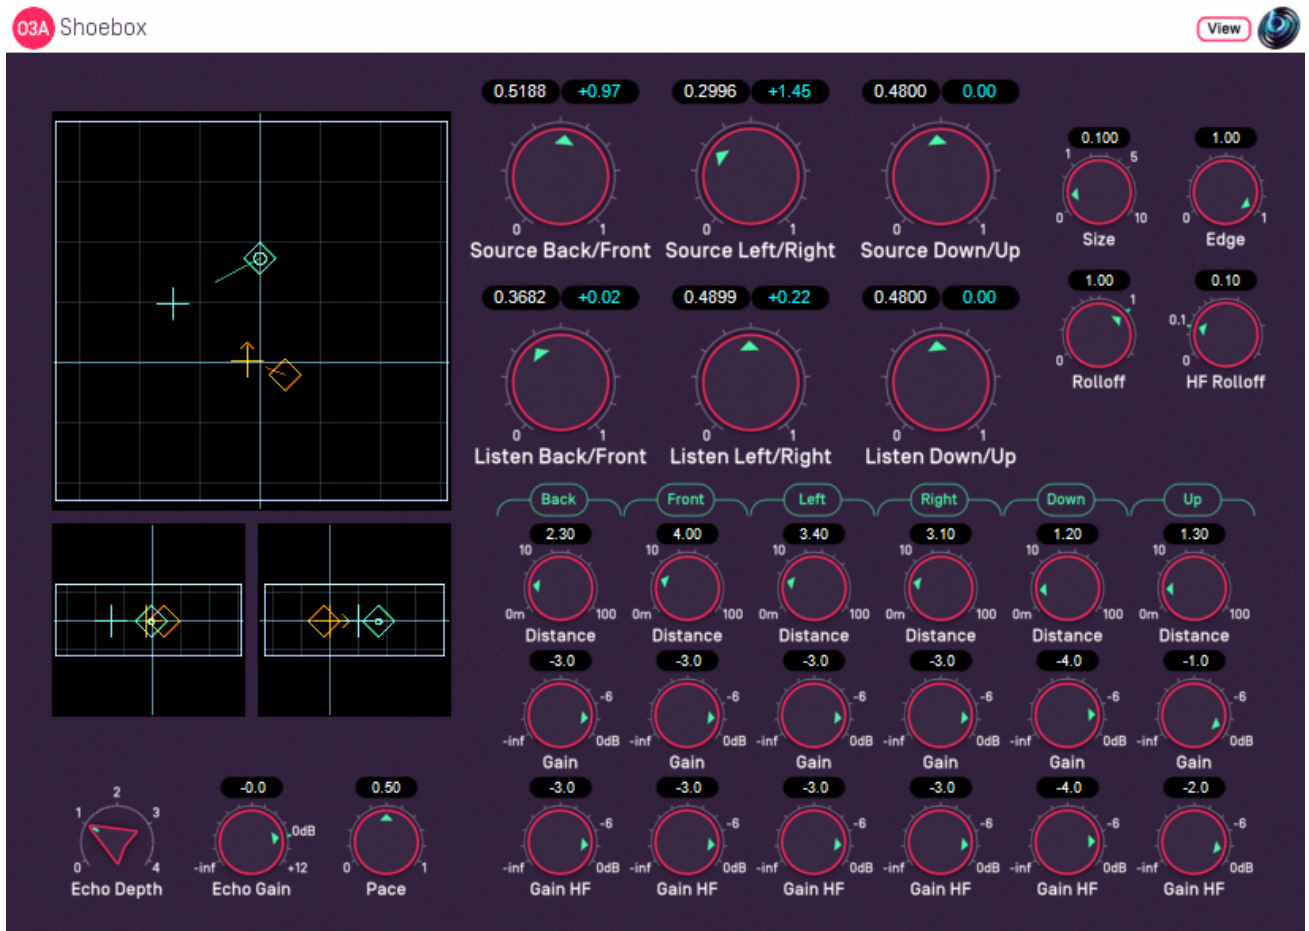
This controls the cutoff frequency of a simple high-pass filter applied to the final output of the late reflection model.

4.3.6 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

4.4 O3A Shoebox



4.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.4.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

4.4.3 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

4.4.4 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.

4.4.4.1 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 O3A Panner LS and O3A Zoom plugins from the O3A 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.

4.4.4.2 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 than 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.

4.4.4.3 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.


4.4.4.4 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 [O3A Rotation](#) or [O3A Look](#) plugins. You can mix the output of several shoebox models together before applying this modification.

4.4.5 View Support

When used with O3A View Sync from the O3A 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.	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.
Hold the left mouse button to set the source 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.

4.4.6 Controls

4.4.6.1 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.)

4.4.6.2 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.)

4.4.6.3 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.

4.4.6.4 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 [O3A 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 O3A Diffuser plugin.

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

4.4.6.5 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.

4.4.6.6 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.

4.4.6.7 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.

4.4.6.8 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. Load is roughly proportional to the number of paths.

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

4.4.6.9 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.

4.4.6.10 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.

4.4.6.11 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.

4.4.6.12 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.


4.4.7 Presets

Presets are available for this plugin:

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

5 O3A Reverb - Convolution Reverb

5.1 O3A Reverb - Convolution



5.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

5.1.3 Controls

- Dry Gain
- Wet Gain
- File Format

5.1.4 Description

This plugin provides basic convolution reverb, but designed to be easy to use with [O3A 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 16 channel O3A input, but this is mostly to allow some of the O3A dry signal to be mixed in to the output. The input to the convolution algorithm itself uses just the first, omnidirectional channel of the O3A 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.

5.1.4.1 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.

For best results, a 16 channel O3A impulse response should be used. Recordings of real spaces done this way are currently relatively rare. However, first or second order ambisonic B-Format impulse responses (4 channel or 9 channel) will work too. They are spatially less precise, but this is often acceptable for 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](#).

5.1.4.2 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 impulse 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/>

5.1.4.3 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 impulse 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.

5.1.4.4 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 O3A 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 O3A file and can be created and/or modified using a variety of O3A plugins. For instance, you might use O3A upmixers to convert a 5.1 impulse response to O3A, 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 [O3A Shoebox](#) plugin. In the latter case, or if you are also using the [O3A 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.

5.1.4.5 Hybrid Reverb

You may also wish to consider using a "hybrid" approach to reverb, for instance using the [O3A 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 [O3A Reverb - Shaped Convolution](#) plugin below.

5.1.4.6 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 O3A 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 O3A 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 O3A input. Basic O3A 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".

5.1.5 Controls

5.1.5.1 Controls: Dry Gain and Wet Gain

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

5.1.5.2 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 and O3A 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.

5.2 O3A Reverb - Shaped Convolution

O3A Reverb - Shaped Convolution

D:\ImpulseResponses\ODEON\HagiaSophia_6.7m_SecondOrder.wav

Envelope Delay	0.030	Envelope Time	0.100
Envelope Start	1.00	Envelope End	0.00
Locality	0.50	Precision	3

D:\ImpulseResponses\ODEON\HagiaSophia_6.7m_SecondOrder.wav
16x16 17.27s 44.1kHz (SRC: 48kHz)
Active.

Dry Gain: -inf to +12 dB

Wet Gain: 3.0, -inf to +12 dB

Pre-Processing Controls

Directivity Shaping Envelope

5.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

5.2.3 Controls

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

5.2.4 Description

This plugin has the features of the simpler [O3A 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 significant CPU loads. Please use it with care.

5.2.4.1 The Directivity Shaping Envelope

The difference between this and the [O3A Reverb - Convolution](#) plugin is the presence of a "directivity shaping envelope" which runs *along the length of the reverb impulse response*. In other words, 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 [O3A 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.

5.2.4.2 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 [O3A Shoebox](#) or [O3A 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 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.

5.2.4.3 Memory Use and CPU Load

This plugin can use even more memory and CPU resources than [O3A 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.

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.

5.2.5 Controls

5.2.5.1 Controls: Dry Gain and Wet Gain

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

5.2.5.2 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.

5.2.5.3 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.

5.2.5.4 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.

5.2.5.5 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.

5.2.5.6 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.

5.2.5.7 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 three indicates that full third order ambisonics (O3A) will be used.

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

5.2.5.8 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 and O3A 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.

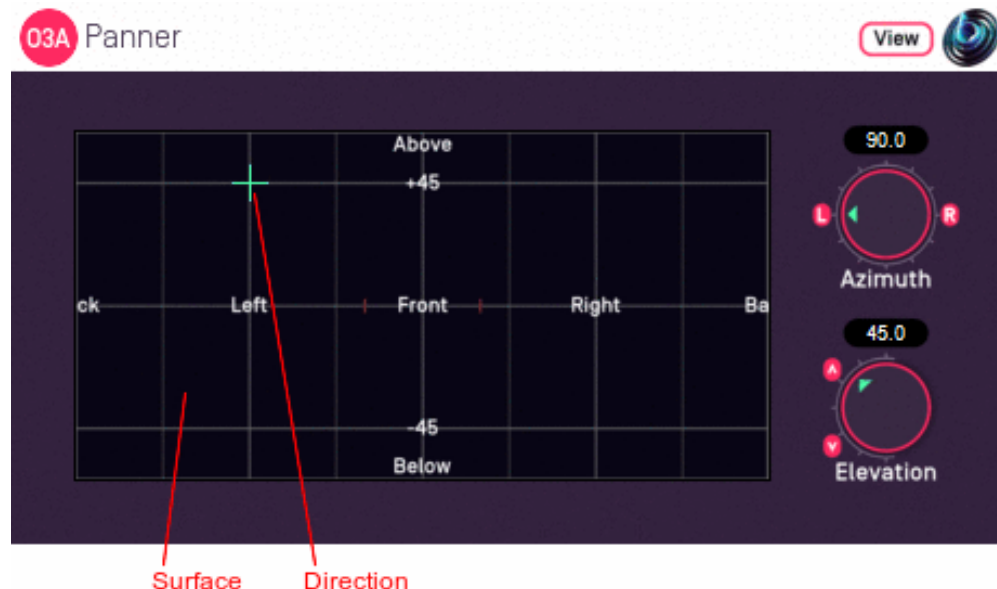
5.2.6 Presets

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

6 O3A Core - Panners

6.1 O3A Panner



6.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.1.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

6.1.3 Controls

- Azimuth
- Elevation

6.1.4 Description

This O3A panner takes a mono sound and places it in a particular direction in a 3D O3A mix.

There are various versions of this panner plugin, but this is our favourite. This version uses more screen real estate than the "classic" version but is typically easier to use, particularly if you are just starting out with this way of doing things.


As well as azimuth and elevation dials, there is a control surface which can be used to visualise the current direction (indicated by the crosshairs). This direction can also be changed by clicking with the mouse.

There are labels on the control surface to help you find left, right, above etc. However, if this still does not make any sense, you might want to see our page on the [O3A Visualiser](#) which lays out its viewing region in the same way.

6.1.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross.	The current direction is shown using a beam and cross.
Hold the left mouse button to set the azimuth and elevation.	Hold a VR controller trigger to set the azimuth and elevation.

6.1.6 Controls

6.1.6.1 Controls: Azimuth and Elevation


Azimuth and Elevation control the direction in which the input is placed in the 3D audio scene.

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.

It may help to experiment with the control surface to get used to this way of doing things.

6.2 O3A Panner - Classic



6.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.2.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

6.2.3 Controls

- Azimuth
- Elevation


6.2.4 Description

There are various versions of this panner plugin and this is the most "raw" version. Only the "classic" azimuth/elevation controls are provided, largely to save screen real estate. These are great if you are used to them; if you are not, you might want to try the [standard version](#).

6.2.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross. Hold the left mouse button to set the azimuth and elevation.	The current direction is shown using a beam and cross. Hold a VR controller trigger to set the azimuth and elevation.

6.2.6 Controls

6.2.6.1 Controls: Azimuth and Elevation


Azimuth and Elevation control the direction in which the input is placed in the 3D audio scene.

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.

It may help to experiment with the [standard version](#) to get used to this way of doing things.

6.3 O3A Panner - Hemisphere



6.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.3.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

6.3.3 Controls

- Azimuth
- Elevation

6.3.4 Description

There are various versions of this panner plugin. This version provides a control surface that maps directly onto the upper hemisphere of possible sound directions. You can click on the control surface to move the sound and holding some keys down changes how the sound is moved:

- Use the "shift" key to select the lower hemisphere (so sounds are below rather than above).
- Use the "alt" key to lock the elevation (so sounds move in a horizontal circle).

A [hemisphere visualiser](#) is available for use with this plugin. However, we generally recommend the [standard panner](#), which has a control surface using the same rectangular layout as the standard

6.3.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross. Hold the left mouse button to set the azimuth and elevation.	The current direction is shown using a beam and cross. Hold a VR controller trigger to set the azimuth and elevation.

6.3.6 Controls

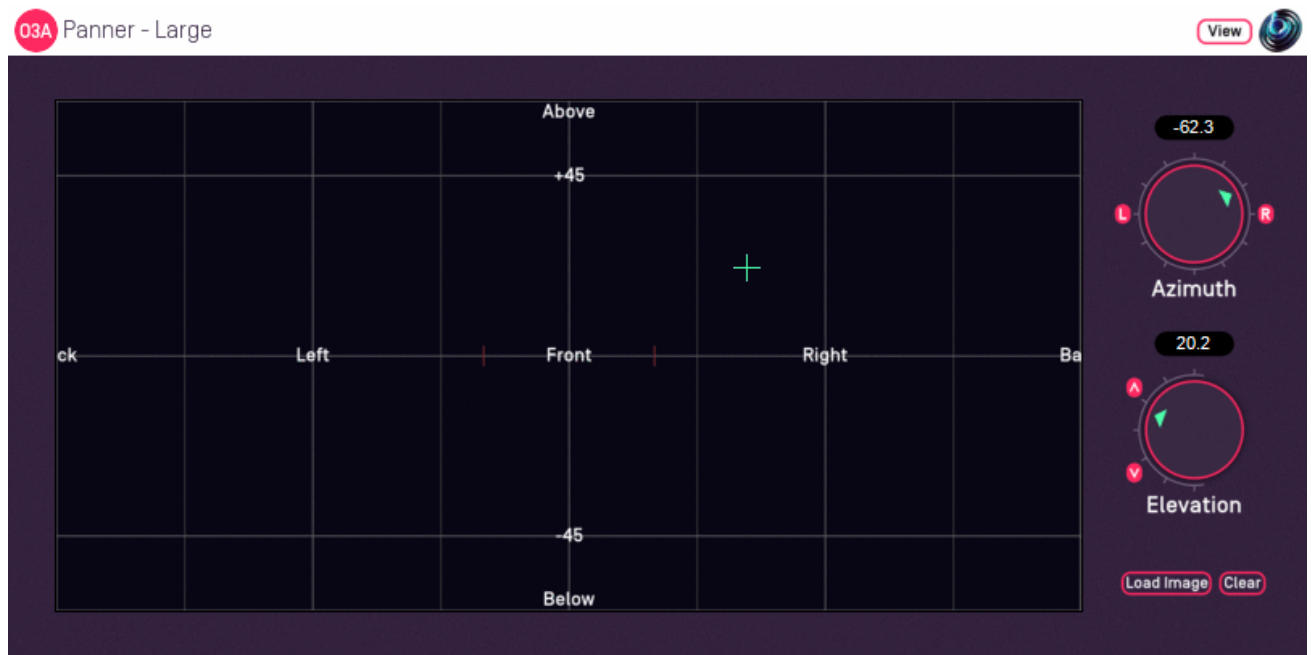
6.3.6.1 Controls: Azimuth and Elevation

Azimuth and Elevation control the direction in which the input is placed in the 3D audio scene.

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.

6.4 O3A Panner - Large



6.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.4.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

6.4.3 Controls

- Azimuth
- Elevation

6.4.4 Description

This O3A panner takes a mono sound and places it in a particular direction in a 3D O3A mix.

This plugin is an extra large version of our [standard panner plugin](#). It can be useful on high resolution screens. It also supports background images.

6.4.4.1 Background Images

Press the "Load Image" button to load an image to replace the default grid background. JPG and PNG images are supported. Because of the way the plugin works, we recommend use of dark backgrounds.


The target image size is 636x316. Other sizes are resized to fit.

The "Clear" button can be used to clear the image so the default grid background is shown again.

6.4.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross. Hold the left mouse button to set the azimuth and elevation.	The current direction is shown using a beam and cross. Hold a VR controller trigger to set the azimuth and elevation.

6.4.6 Controls

6.4.6.1 Controls: Azimuth and Elevation


Azimuth and Elevation control the direction in which the input is placed in the 3D audio scene.

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.

It may help to experiment with the control surface to get used to this way of doing things.

6.5 O3A Panner - Two Channel



6.5.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.5.2 Audio

	Channels	Content
Input	2	Stereo
Output	16	O3A

6.5.3 Controls

- Azimuth 1
- Elevation 1
- Azimuth 2
- Elevation 2
- Mouse Mode

6.5.4 Description

This O3A panner takes two mono sound sources and places them in particular directions in a 3D O3A mix.

It can also be used as a simple way to upmix stereo to O3A. However, we do not particularly

recommend you do this except for fairly narrow angles because a large "hole" can appear between the two sources. For a more thorough handling of material that was originally panned to stereo you wish to look at the O3A Upmixer - Stereo plugin.

There are azimuth and elevation dials for each of the two sources. Each works the same way as the controls in the [mono version](#). There is a control surface which can be used to visualise the current directions for the two sources, indicated by the green/blue and red crosshairs. These can be changed by clicking with the mouse, using the left mouse button for the first source and the right mouse button for the second. If you do not have a right mouse button, or are using a touch screen, you can hold down the 'Alt' key and press the main button instead. You can also move both sources around together using the 'Shift' key.

If pressing keys on the keyboard is inconvenient, there is a mouse mode slider on screen which can be used instead.


Mouse Mode	Key(s) Held	Function
Norm		Normal operation. Left click to move the first source, right click to move the second source.
Alt	Alt	Flips the source moved when the control surface is clicked, so a left click moves the second source rather than the first (and vice versa). Typically used to select the second source on a mouse with only one button.
Both	Shift	Changes the control surface so that movements of the mouse shift both sources.
Opp	Alt & Shift	Changes the control surface so that movements of the mouse shift both sources, but the sources move in opposite directions.

There are labels on the control surface to help you find left, right, above etc. However, if this still does not make any sense, you might want to see our page on the [O3A Visualiser](#) which lays out its viewing region in the same way.

6.5.5 View Support

When used with O3A View Sync from the O3A 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
<p>The current panner directions are shown using green/blue and red crosses, similarly to the plugin display.</p> <p>Point at one and hold the left mouse button to select it. Then, drag to set its azimuth and elevation.</p>	<p>The current panner directions are shown using green/blue and red beams and crosses, similarly to the plugin display.</p> <p>Each VR controller moves one panner direction and the controllers show short green/blue and red beams to indicate which is which. Hold a VR controller trigger to set an azimuth and elevation.</p>

6.5.6 Controls

6.5.6.1 Controls: Azimuth 1 and 2 and Elevation 1 and 2

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.



Azimuth 1 and Elevation 1 control the location used to pan the first source (which is shown in green/blue on the control surface). Azimuth 2 and Elevation 2 control the location used to pan the first source (which is shown in red).

6.5.6.2 Control: Mouse Mode

The mouse mode slider in the middle at the bottom of the screen determines the mouse mode used when the control surface is clicked. They are an alternative to use of the Alt and Shift keys. See the [table above](#) for details on how these options affect operation.

6.6 O3A Panner - Eight Channel

O3A Panner - Eight Channel



6.6.1 Host Support

Host Type	Support
AAX	Yes, input mapped to 7.1
VST2	Yes

6.6.2 Audio

	Channels	Content
Input	8	Eight channels of mono (mapped to 7.1 for AAX)
Output	16	O3A

6.6.3 Controls

- Azimuth 1 - 8
- Elevation 1 - 8
- Gain 1 - 8

- Mute 1 - 8
- Solo 1 - 8

6.6.4 Description

This O3A panner takes eight mono sound sources and places them in particular directions in a 3D O3A mix.

There are azimuth and elevation dials for each of the sources. These work the same way as the controls in the standard [single channel panner](#). Each source also has mute and solo buttons and a gain control.

There is a control surface which can be used to visualise the current directions for the sources, indicated by coloured crosshairs. These can be changed by dragging with the mouse.

There are labels on the control surface to help you find left, right, above etc. However, if this still does not make any sense, you might want to see our page on the [O3A Visualiser](#) which lays out its viewing region in the same way.

Text can be entered into channel labels above the mute/solo buttons to help you keep track of which channel is which. these labels are shown when used with this plugin's [View support](#).

6.6.4.1 Reaper Note

If you are using this panner with Reaper, it is fairly easy to set up the individual mono source inputs as subtracks of the track on which the panner is set.

On the routing panel for each subtrack, set the "Parent Channels" to send to the relevant pair of inputs in this plugin and then pan hard left or right depending on whether this should be an odd or even input. Set this panner plugin as the first effect on the parent track and check the routing using the solo button for each source.

6.6.4.2 Pro Tools Note

The AAX version of this plugin expects input as 7.1 audio, with the eight panning channels mapped from the 7.1 channels in the following order:


Index	Label
1	L
2	C
3	R
4	Lss
5	Rss
6	Lsr
7	Rsr
8	LFE

A simple way to use this plugin is to place it as the first plugin on a 7.1 track. Mono signals can then be sent from other tracks by addressing the individual channels of the track's 7.1 input bus. For instance, if the input bus is labelled "Panner8" then to send to the first panner channel, mono audio should be sent to "Panner8.L". To send to the second panner channel, "Panner8.C" and so on, through to "Panner8.LFE" for the eighth channel. Routing can be checked using the plugin's solo buttons in the usual way.

6.6.5 View Support

When used with O3A View Sync from the O3A 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
<p>The current panner directions are shown using colour-coded crosses and labels, similarly to the plugin display.</p> <p>Point at one and hold the left mouse button to select it. Then, drag to set its azimuth and elevation.</p>	<p>The current panner directions are shown using colour-coded beams, crosses and labels, similarly to the plugin display.</p> <p>Point at one with a VR controller and the controller will indicate the label of the source pointed at. Then, hold the VR controller trigger to select it and move the controller to set its azimuth and elevation.</p>

6.6.6 Controls

6.6.6.1 Controls: Azimuth 1 - 8 and Elevation 1 - 8

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.

It may help to experiment with the control surface to get used to this way of doing things.

Azimuth and Elevation can also be changed by dragging the relevant source on the control surface.

6.6.6.2 Controls: Gain 1 - 8

Each panned source has its own gain control, with a range from -60dB to +12dB.

6.6.6.3 Controls: Mute 1 - 8 and Solo 1 - 8

Each panned source has its own mute and solo switch. These have the usual meanings and solo overrides mute.

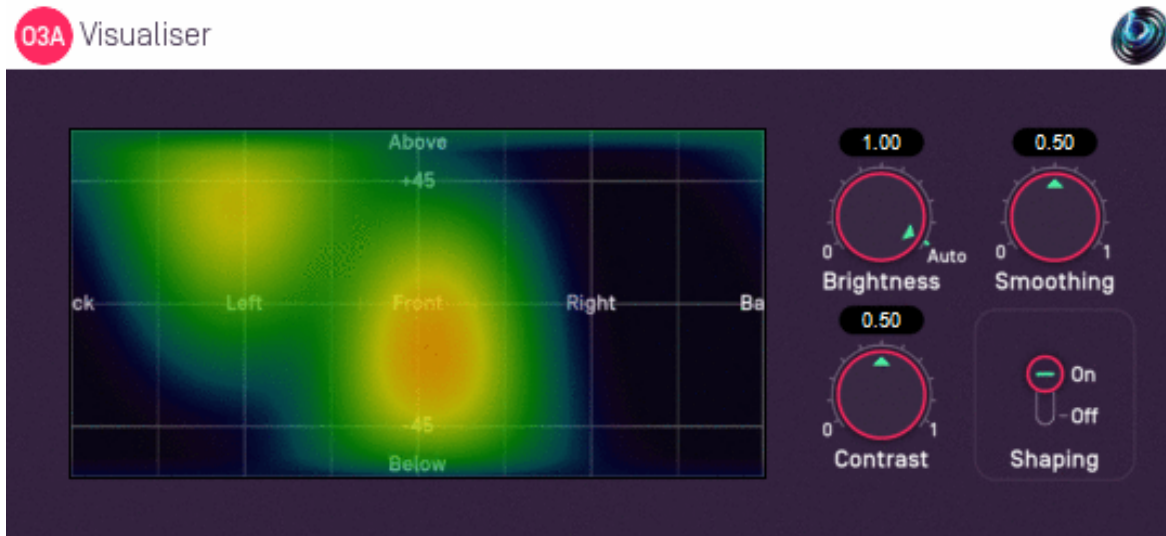
6.6.7 Presets

Presets are available for this plugin:

- 5.0
- 7.0
- Triangle
- Square
- Pentagon
- Hexagon
- Octagon
- Front Stage
- Bi-Triangle
- Octahedon
- Cube

7 O3A Core - Visualisers

7.1 O3A Visualiser



7.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

7.1.3 Controls

- Brightness
- Contrast
- Smoothing
- Shaping

7.1.4 Description

This plugin produces a view of an O3A stream which can help make sense of what is happening in it. Audio is passed through unmodified, so the plugin can be inserted in the middle of an O3A effects chain.

The stream is shown using a rectangular screen region. This rectangular region shows an equal-area cylindrical projection of the directional components of the O3A soundfield interpreted over a sphere.

Or, to put it another way, it lights up in different places depending on where a sound comes from. Some of those places are labelled in the image. The view is like a map of the world, where sounds above are at the "north pole" and are shown across the top of the view and sounds below are at the

"south pole" at the bottom. Sounds rotating around at head height move around the equator, so they move along the view halfway down, wrapping around between the left and right edges as they pass behind.

A good way to try this out is to chain the [O3A Panner](#) immediately before this plugin as this uses the same directional approach. Then set up a decoder after this suitable for use with your speakers or headphones. Then you can play around with different sound directions and hear *and see* the results.

To reduce CPU load, this visualiser stops most processing when it is not on screen.

This plugin does not have O3A View support. You might wish to consider [O3A Flare](#).

7.1.5 Controls

7.1.5.1 Control: Brightness

This controls the brightness of the image. If the screen is black when sounds are playing, you may wish to turn this up. If it is all red, turn it down.

Alternatively, you can set the dial to "Auto", in which case the plugin will try to ride the brightness control to keep the visualisation useful.

7.1.5.2 Control: Contrast

This controls the contrast of the image. Higher values produce a sharper image but map a smaller dynamic range onto the screen.


7.1.5.3 Control: Smoothing

The image generated is smoothed over time. Use more smoothing for a more stable, slower moving image. Use less for a more responsive one.

7.1.5.4 Control: Shaping

Shaping should normally be left on as it produces a more straightforward view. If this is turned off, the image will typically be sharper, but ripples will be visible in directions away from original sound directions.

7.2 O3A Visualiser - Hemisphere



7.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

7.2.3 Controls

- Brightness
- Contrast
- Smoothing
- Shaping
- Viewpoint

7.2.4 Description

This plugin produces a view of an O3A stream which can help make sense of what is happening in it. Audio is passed through unmodified, so the plugin can be inserted in the middle of an O3A effects chain.

The stream is shown using a two circular regions each showing a projection of a half of the directional components of the O3A soundfield interpreted over a sphere.

By default, the left sphere shows the upper hemisphere of the soundfield and the right sphere shows the lower hemisphere, as if both were viewed from above. The image is orientated so that sounds at the front appear at the top of both images. However, the [viewpoint](#) can be changed.

A good way to try this out is to chain the [O3A Panner - Hemisphere](#) plugin immediately before this one as that panner uses the same directional approach (as long as the [viewpoint](#) is set to "Above"). Then set up a decoder after this suitable for use with your speakers or headphones. Then you can play around with different sound directions and hear and see the results.

We do not particularly recommend this visualisation. Generally, we suggest you use the standard [O3A Visualiser](#) instead.

This plugin does not have O3A View support. You might wish to consider [O3A Flare](#).

7.2.5 Controls

7.2.5.1 Control: Brightness

This controls the brightness of the image. If the screen shows nothing when sounds are playing, you may wish to turn this up. If it is all red, turn it down.

Alternatively, you can set the dial to "Auto", in which case the plugin will try to ride the brightness control to keep the visualisation useful.

7.2.5.2 Control: Contrast

This controls the contrast of the image. Higher values produce a sharper image but map a smaller dynamic range onto the screen.

7.2.5.3 Control: Smoothing

The image generated is smoothed over time. Use more smoothing for a more stable, slower moving image. Use less for a more responsive one.

7.2.5.4 Control: Shaping


Shaping should normally be left on as it produces a more straightforward view. If this is turned off, the image will typically be sharper, but ripples will be visible in directions away from original sound directions.

7.2.5.5 Control: Viewpoint

By default, the visualisation shows the upper hemisphere on the left and the lower hemisphere on the right, both viewed as the viewpoint was from **above**.

Alternatively, the viewpoint can be changed to view from the **right** or from **behind**. In both cases sounds above appear at the top of the image.

7.3 O3A Colourizer



7.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.3.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

7.3.3 Controls

- Hue Cutoff
- Smooth Cutoff
- Rider Cutoff
- Brightness

7.3.4 Description

This plugin is a variant of the [O3A Visualiser](#) plugin. Like that plugin, it produces a view of an O3A stream which can help understand what is happening in it. Audio is passed through unmodified, so the plugin can be inserted in the middle of an O3A effects chain.

Two significant differences between this and the [O3A Visualiser](#) are the size of the plugin and how colours are mapped. With this plugin, colours are chosen using a hue that depends on the frequency content of the sound. Sounds that are predominantly low-frequency are coloured red and sounds that are predominantly high-frequency are coloured magenta. Sounds between are coloured with other hues accordingly.

7.3.4.1 Background Images

Press the "Load Image" button to load an image to replace the default grid background. JPG and PNG images are supported. Because of the way the plugin works, we recommend use of dark backgrounds.

The target image size is 636x316. Other sizes are resized to fit.

The "Clear" button can be used to clear the image so the default grid background is shown again.

7.3.5 Controls

7.3.5.1 Control: Hue Cutoff

The plugin selects hues by splitting the sound into low and high frequencies. To do this, it uses a low-pass filter with a cutoff frequency given by this control. Generally, it is best to set this where you want the greatest sensitivity to frequency content.

7.3.5.2 Control: Smooth Cutoff

The local image level is smoothed over time using a low-pass filter with a cutoff frequency given by this control. Lower values will give a more stable, slower moving image. Higher values will give a more responsive one.


7.3.5.3 Control: Rider Cutoff

A "gain rider" tracks the overall signal level using another low-pass filter. Lower values respond more slowly to changes in the overall level of the signal. Higher values respond more quickly.

7.3.5.4 Control: Brightness

This modifies the overall brightness of the image.

7.4 O3A Flare



7.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.4.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

7.4.3 Controls

- Rider Cutoff
- Brightness

7.4.4 Description

This plugin is a variant of the [O3A Visualiser](#) plugin. Like that plugin, it produces a view of an O3A stream which can help understand what is happening in it. Audio is passed through unmodified, so the plugin can be inserted in the middle of an O3A effects chain.

Although this plugin uses the same layout, it uses substantially different processing to the [O3A Visualiser](#) and [O3A Colourizer](#) plugins. This plugin uses simple first order estimates of the predominant sound direction in a range of frequency bands. These directions are painted using colour coding, ranging from red for low frequencies up to magenta for high frequencies (up to about 20kHz). Plot accuracy is better for simpler scenes; multiple sources with similar frequency content will often result in smudged or misaligned images.

To reduce CPU load, this visualiser stops most processing when it is not on screen.

7.4.4.1 Background Images

Press the "Load Image" button to load an image to replace the default grid background. JPG and PNG images are supported. Because of the way the plugin works, we recommend use of dark backgrounds.


The target image size is 636x316. Other sizes are resized to fit.

The "Clear" button can be used to clear the image so the default grid background is shown again.

7.4.5 View Support

When used with O3A View Sync from the O3A 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 flare visualisation is supported by View.	The flare visualisation is supported by ViewVR.

Warning: requires a fast network connection and powerful CPU.

7.4.6 Controls

7.4.6.1 Control: Rider Cutoff

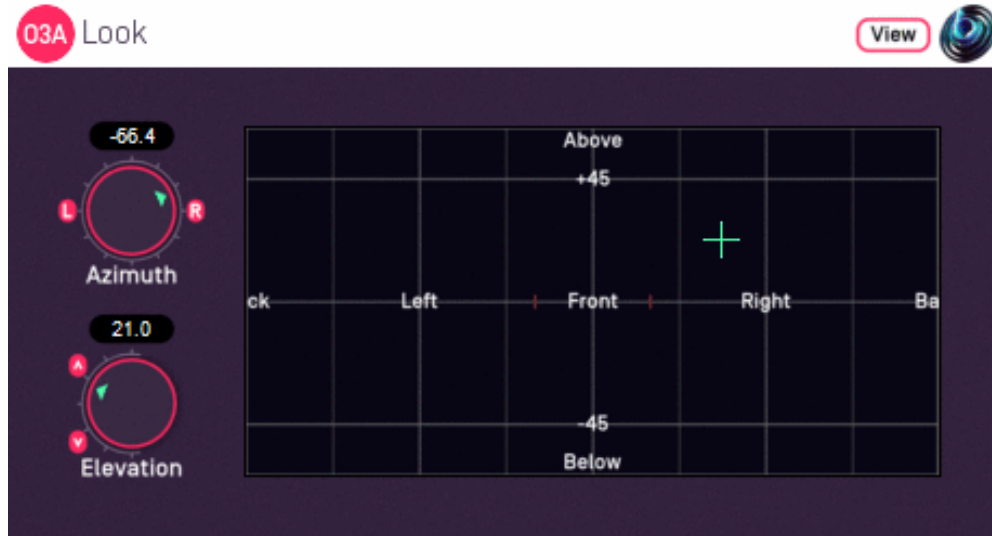
A "gain rider" tracks the overall level using another low-pass filter. Lower values respond more slowly to changes in the overall level of the signal. Higher values respond more quickly.

7.4.6.2 Control: Brightness

This modifies the overall brightness of the image.

8 O3A Core - Rotation

8.1 O3A Look



8.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

8.1.3 Controls

- Azimuth
- Elevation

8.1.4 Description

This plugin changes where "front" is in the 3D audio scene, as if the listener had turned their head to look (and listen) in another direction.


The "look" direction is chosen using the dials or control surface, in the normal way (for instance, see the [O3A Panner](#)). Sounds that were in this direction will be to the front after processing.

This is essentially an simplified (and inverted) version of the [O3A Rotation](#) plugin. The "roll" and "ordering" controls are not provided. Instead, we just assume that the listener turns their head and pitches it backwards or forwards. Rolling to the sides is not supported (use the rotation plugin if you need this).

8.1.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross. Hold the left mouse button to set the azimuth and elevation.	The current direction is shown using a beam and cross. Hold a VR controller trigger to set the azimuth and elevation.

8.1.6 Controls

8.1.6.1 Controls: Azimuth and Elevation

Azimuth and Elevation control the direction that is brought to the front by this plugin.

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.

8.2 O3A Rotation



8.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

8.2.3 Controls

- Yaw
- Pitch
- Roll
- Ordering

8.2.4 Description

This plugin rotates the entire O3A soundfield around the centre of the room.

Any rotation around the centre is possible and is built up using yaw, pitch and roll rotations. You can think of the rotations as being applied in order.


We *strongly recommend* that you use the [O3A Visualiser](#) when working with this plugin, particularly if you are going to use more than one of the dials at once. Otherwise, this plugin can be very confusing! It is easy to control in Virtual Reality using the ViewVR application however.

If you are using this plugin to simulate changes to the listener's orientation, you might also want to consider the [O3A Look](#) plugin, which is simpler to use but does not include a Roll control.

8.2.5 View Support

When used with O3A View Sync from the O3A 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
<p>The current post-rotation 'front' of the scene is shown using a cross.</p> <p>Hold the left mouse button to set all four rotation controls, moving the 'front' of the scene to the indicated direction. Only 2DOF is supported and the rotation will be set so 'up' is upwards.</p>	<p>The scene orientation is shown using a floating cube with matching marker beams on the cube and VR controllers.</p> <p>Rotate a VR controller into the chosen orientation and hold the trigger to set the four rotation controls. Full 3DOF is supported.</p>

8.2.6 Controls

8.2.6.1 Controls: Pitch, Roll and Yaw

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.

This plugin can be used to simulate changes in the listener's orientation. If you are doing this, you will probably want to change the sign of the angles involved. For instance, turning the listener's head to the right means that the listener hears sound move to the left.

8.2.6.2 Control: Ordering


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).

When simulating changes in the listener's orientation, we generally recommend the Y-P-R ordering. Then, the three dials can be thought of as turning the listener's head horizontally, then pitching it

forwards or backwards and finally rolling it from side to side.

9 O3A Core - Decoders and Converters

9.1 O3A Decoder - Stereo



9.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.1.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Stereo

9.1.3 Controls

- [Method](#)

9.1.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes (renders) it as a simple, robust stereo mix.

Please note that **this is not a binaural decoder**. You may wish to consider [O3A Decoder - Headphones](#).

As another alternative to this plugin, you can use the [O3A Virtual Microphone](#) plugin to produce a stereo mix using a simple virtual stereo microphone.

9.1.5 Controls

9.1.5.1 Control: Method

Two decoder methods are supported, "Panner" and "Basic Head".

Method	Description
Panner	<p>This method produces a simple stereo mix, largely as if the sounds within the mix had been panned conventionally. This can work well on speakers or headphones.</p> <p>Sounds from the front and back are folded down on the stereo image using different angle conventions. Specifically, sounds panned from -30 to +30 degrees at the front will pass across the stereo image, as will sounds panned from -110 to +110 degrees at the back. Sounds above or below are fed to both speakers.</p>
Basic Head	<p>This method produces a stereo mix using a very crude model of a human head. When listened to on headphones, this imitates typical left and right ear levels that would occur naturally, albeit not in a frequency-dependent way. This can work well on speakers or headphones.</p> <p>This is not a binaural decoder. Although the data for this method is based on the Amber HRTF, it has been simplified so that the output is not frequency-dependent, for better results on loudspeakers. This reduces the strength of 3D cues on headphones.</p> <p>For full binaural decoding, we recommend O3A Decoder - Headphones from the O3A Decoding library, O3A Decoder - VR Monitoring from the O3A View plugin library, or Rapture3D Advanced.</p> <p>Unlike most other O3A decoders, this decoder presents sounds from different directions at different levels. For instance, sounds to the front are louder than those behind.</p>

9.1.5.2 Technical Notes

The Blue Ripple Sound Amber HRTF uses data from the IRCAM LISTEN HRTF data set, available at <http://recherche.ircam.fr/equipes/salles/listen/index.html>. It has been processed and simplified heavily.

9.2 O3A Decoder - Mono

 O3A Decoder - Mono



9.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.2.2 Audio

	Channels	Content
Input	16	O3A
Output	1	Mono

9.2.3 Description

This plugin takes a third order ambisonics (O3A) mix and reduces it to mono by discarding everything but its omnidirectional component.

This type of mix is particularly suitable for level analysis, including loudness monitoring.

If you wish to extract a mono response from the soundfield in a particular direction, you might also want to try the [O3A Virtual Microphone](#) or [O3A Beamer](#) plugins.

9.3 O3A Decoder - 5.1 Basic



Decoder - 5.1 Basic



9.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.3.2 Audio

	Channels	Content
Input	16	O3A
Output	6	5.1 Surround

9.3.3 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as a 5.1 surround mix.

A resilient decoding is used, suitable for small or large venues. ITU speaker angles are assumed.

A version of this plugin with more controls is available in the O3A Decoding plugin library.


9.3.3.1 Channels

The channel ordering used is:

Channel	AAX	VST
1	Front Left	Front Left
2	Front Centre	Front Right
3	Front Right	Front Centre
4	Left Surround	LFE
5	Right Surround	Side Left
6	LFE	Side Right

The low frequency effect channel is *not* used.

9.4 O3A Decoder - FuMa



9.4.1 Host Support

Host Type	Support
AAX	Yes, output mapped to standard SN3D ambisonics
VST2	Yes

9.4.2 Audio

	Channels	Content
Input	16	O3A
Output	16	FuMa Ambisonic Audio (mapped to standard SN3D ambisonics for AAX)

9.4.3 Description

This plugin takes a third order ambisonic mix prepared with these plugins (O3A) and converts it to a third order mix using the FuMa ambisonic convention. The FuMa convention was used in early versions of these plugins (prior to version 2.0) and is an extension of classic First Order Ambisonics.

We have labelled this plugin as a decoder hopefully to make it clear that its output is an external format and no longer compatible with the [convention](#) that the rest of these plugins use.

You can convert back using the [O3A Injector - FuMa](#) plugin.

In Pro Tools, as the FuMa format is not supported directly as a stem format, this plugin outputs audio using the ordinary ambisonic stem format. Normally this should be encoded as O3A (i.e. SN3D/ACN), so we are "misusing" the format here.

9.4.3.1 Orders

This plugin produces third order FuMa output. If you only want first or second order output, just take the first four or nine output channels and ignore the rest.

Alternatively, the [O3A Decoder - O1A](#) or [O2A](#) plugins can be "misused" to reduce the order of the FuMa material as this is a rare case where a plugin will work correctly for both SN3D and FuMa.

9.5 O3A Injector - FuMa



Injector - FuMa



FuMa 3 -> O3A [SN3D, AmbiX]

9.5.1 Host Support

Host Type	Support
AAX	Yes, input mapped to standard SN3D ambisonics
VST2	Yes

9.5.2 Audio

	Channels	Content
Input	16	FuMa Ambisonic Audio (mapped to standard SN3D ambisonics for AAX)
Output	16	O3A

9.5.3 Description

This plugin takes a mix in the FuMa ambisonic convention and converts it to [third order ambisonics](#) (O3A) as used by the current version of these plugins. FuMa was used by old versions of these plugins (prior to version 2.0) and is an extension of classic First Order Ambisonics.

You can convert back with the [O3A Decoder - FuMa](#) plugin.

In Pro Tools, as the FuMa format is not supported directly as a stem format, this plugin expects audio using the ordinary ambisonic stem format. Normally this should be encoded as O3A (i.e. SN3D/ACN), so we are "misusing" the format here.

9.5.3.1 Orders

This plugin expects third order FuMa input. If you only have first or second order, feed this to the first four or nine input channels and feed silence to the rest.

Alternatively, the [O3A Injector - O1A or O2A plugins](#) can be "misused" to increase the order of the FuMa material as this is a rare case where a plugin will work correctly for both SN3D and FuMa.

9.6 O3A Decoder - O1A and O2A

 O3A Decoder - O1A



O3A [SN3D, AmbiX] -> O1A [SN3D, AmbiX]

9.6.1 Host Support

9.6.1.1 O3A Decoder - O1A

Host Type	Support
AAX	Yes
VST2	Yes

9.6.1.2 O3A Decoder - O2A

Host Type	Support
AAX	Yes
VST2	Yes

9.6.2 Audio

9.6.2.1 O3A Decoder - O1A

	Channels	Content
Input	16	O3A
Output	4	O1A

9.6.2.2 O3A Decoder - O2A

	Channels	Content
Input	16	O3A
Output	9	O2A

9.6.3 Description

These simple plugins take a third order (O3A) ambisonic mix and convert it to a lower resolution first order (O1A) or second order (O2A) mix. In all cases, the SN3D convention is used.

These plugins are extremely simple and just reduce the channel count in use by discarding the later channels, which has the effect of reducing the spatial detail of the mix.

Note that some DAWs (such as Reaper) pass through channels that are not written to. These plugins will appear to have no effect in these DAWs!

These plugins can also be "misused" to reduce the order of N3D/ACN or FuMa material. This is a rare case where the same plugin will work correctly for all three ambisonic formats.

9.7 O3A Injector - O1A and O2A

 O3A Injector - O1A



O1A [SN3D, AmbiX] -> O3A [SN3D, AmbiX]
Warning: output has limited spatial detail.

9.7.1 Host Support

9.7.1.1 O3A Injector - O1A

Host Type	Support
AAX	Yes
VST2	Yes

9.7.1.2 O3A Injector - O2A

Host Type	Support
AAX	Yes
VST2	Yes

9.7.2 Audio

9.7.2.1 O3A Injector - O1A

	Channels	Content
Input	4	O1A
Output	16	O3A

9.7.2.2 O3A Injector - O2A

	Channels	Content
Input	9	O2A
Output	16	O3A

9.7.3 Description

These simple plugins take a lower resolution first order (O1A) or second order (O2A) mix and convert it to a third order (O3A) ambisonic mix. In all cases, the SN3D convention is used.

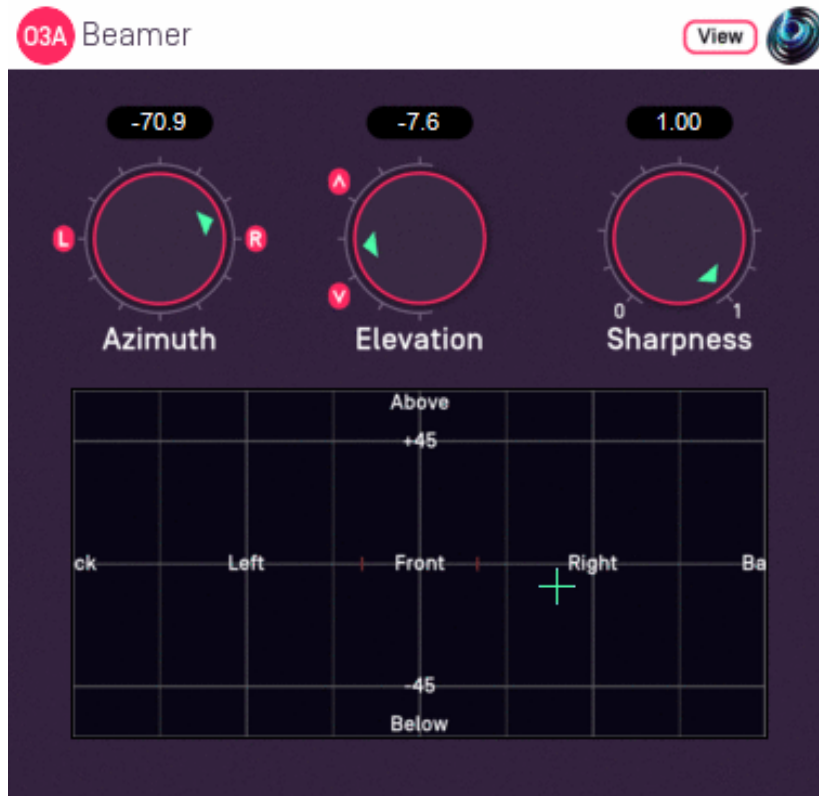
These plugins are extremely simple and just increase the channel count in use, filling new channels with silence. Although this increases the order of the mix, it does not increase its spatial detail. Further, the resulting mix may have unnatural features acoustically.

More sophisticated means to convert first order (O1A) to third order include the O3A Injector - First Order plugin from the O3A Upmixers pack and the O3A Harpex Upsampler.

These plugins can also be "misused" to increase the order of N3D/ACN or FuMa material. This is a rare case where the same plugin will work correctly for all three ambisonic formats.

10 O3A Core - Virtual Microphones

10.1 O3A Beamer



10.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

10.1.2 Audio

	Channels	Content
Input	16	O3A
Output	1	Mono

10.1.3 Controls

- Azimuth
- Elevation
- Sharpness

10.1.4 Description

The O3A Beamer extracts sound in a particular "beam" direction from a 3D O3A mix. You can think of this as a narrow directional microphone response, pointed in a particular direction in the soundfield.

The beam direction can be set using the azimuth and elevation controls, or the control surface, where the direction is indicated by the crosshairs.

There are labels on the control surface to help you find left, right, above etc. However, if this still does not make any sense, you might want to see our page on the [O3A Visualiser](#) which lays out its viewing region in the same way.


The algorithm used is a "passive" one and should not colour frequency content significantly.

If you are interested in producing simple virtual microphone responses from an O3A mix, you may also wish to read about the [O3A Virtual Microphone](#).

10.1.5 View Support

When used with O3A View Sync from the O3A 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 direction is shown using a cross.	The current direction is shown using a beam and cross.
Hold the left mouse button to set the azimuth and elevation.	Hold a VR controller trigger to set the azimuth and elevation.

10.1.6 Controls

10.1.6.1 Controls: Azimuth and Elevation

Azimuth and Elevation control the direction to extracted.


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.

10.1.6.2 Control: Sharpness

The sharpness control determines how directional the beam response is. A value of zero reduces the response to an omnidirectional one like that of the [mono decoder](#). A value of one is the usual case and gives the sharpest response.

10.2 O3A Virtual Microphone



10.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

10.2.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Stereo

10.2.3 Controls

- Stereo Width
- Response

10.2.4 Description

The O3A Virtual Microphone extracts a stereo image from a 3D O3A mix using a front-facing "virtual" stereo microphone placed at the centre of the soundfield. The angle between the two capsules of this virtual microphone can be varied, along with their directional response. Available directional responses include omni, cardioid and figure-of-eight. This makes it possible to synthesise the results of various conventional coincident stereo recording techniques such as Blumlein X-Y and M/S.

This plugin does not allow you to point the front of the virtual microphone in other directions. If the front is not where you want, remember that you can rotate the soundfield using [O3A Rotation](#) or [O3A Look](#).

This plugin actually only needs first order ambisonic information (and even then, the height channel is not used).

Be careful using this as a stereo decoder, because the overall output level can be highly dependent on source locations, as with the "real" recording techniques. You might wish to consider [O3A Decoder - Stereo](#).

If you are interested in extracting sounds in particular directions from an O3A mix, you may also be interested in the sharper response of the [O3A Beamer](#) plugin.

10.2.5 Controls

10.2.5.1 Control: Stereo Width

This is the angle between the two virtual microphones, between 0 and 180 degrees. For instance, 90 degrees should be used for Blumlein X-Y.

10.2.5.2 Control: Response

The response control determines the directional response of the two microphone capsules. Some useful values are:

Response Value	Name
0.00	Omnidirectional
0.25	Subcardioid
0.50	Cardioid
0.75	Hypercardioid
1.00	Figure-of-Eight

For instance, 1.00 (figure-of-eight) should be used for Blumlein X-Y.


10.2.6 Presets

Presets are available for this plugin:

- Cardioids 130
- Cardioids 90
- Blumlein X/Y

11 O3A Core - Metering

11.1 O3A Meter - Signal



11.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

11.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

11.1.3 Description


This extremely simple plugin measures the signal peak in each of the 16 channels of an [O3A stream](#) to indicate if a signal is present. It is intended for basic diagnostics (e.g. "are all my channels getting through?") and not for serious signal analysis. Audio is passed through unchanged.

11.1.3.1 Meters

Lights are shown for each of the 16 channels in O3A. They are shown in rows corresponding to the ambisonic orders. Lights switch on with a yellow colour in the presence of very low signal levels (-100dB and above) and switch to red at high levels (0dB and above). These signal levels are simple peak levels, *not* "True Peak".

Note that first order ambisonic B-Format only uses the first four channels of O3A, in which case only the first four lights (top two rows) will be on. Sounds exactly on the horizontal plane switch on alternating lights on each row and a few special directions do other strange things. However, normal 3D streams should switch on all 16 lights; if they do not there is probably something wrong! A common error is to forget to set a track's channel count or send to at least 16, and this plugin will usually make that obvious.

11.2 O3A Meter



11.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

11.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

11.2.3 Description

This plugin measures the sound level in an O3A stream. A breakdown by ambisonic order is provided, along with a history graph and the current overall level. Audio is passed through unchanged.

11.2.3.1 Meters

From left to right, the following displays are shown:

1. Bar meter showing the current and recent peak level in order 0 (ACN channel 0 only).
2. Bar meter showing the current and recent peak level in order 1 (ACN channels 1 to 3).
3. Bar meter showing the current and recent peak level in order 2 (ACN channels 4 to 8).
4. Bar meter showing the current and recent peak level in order 3 (ACN channels 9 to 15).
5. History min/max graph of overall level (all 16 channels).
6. Bar meter showing the overall current and peak level (all 16 channels).

Each bar meter shows the current RMS level in green, averaged using a 400ms rectangular window and translated to a decibel scale, along with a recent peak calculation using a 2s memory. The history graph shows the last three minutes of the overall level, displaying the range from minimum to maximum for each second.


The meters are all normalised so that a 0dBFS peak sine wave panned in any direction will produce a -3.01dB output. There is no sensitivity to frequency or direction. If you are interested in measuring loudness, you may wish to consider the O3A Meter - Karma plugin from the O3A Decoding

Decoding library, which will produce comparable numbers when set to use LUFS units, subject to directional and loudness weighting.

For material that has been assembled by simple panning, the bar meters should be roughly in line with each other.

12 O3A Core - Soundfield Modification

12.1 O3A Gain



12.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

12.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

12.1.3 Controls

- Gain

12.1.4 Description


This trivial plugin applies a gain to the O3A stream. All channels of the stream are affected equally.

12.1.5 Controls

12.1.5.1 Control: Gain

The gain to apply, in decibels, between -30dB and +30dB.

12.2 O3A Spatial Delay



12.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

12.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

12.2.3 Controls

- Delay Type
- Delay Time
- Dry Gain
- Wet Gain
- Feedback Gain
- LPF Cutoff
- Reflection
- Yaw
- Pitch
- Roll
- Ordering
- Clear Delay

12.2.4 Description

This plugin delays a third order ambisonic (O3A) stream, applying a scene reflection, rotation and low pass filter to the delayed sound. It also supports feedback, so the output of the plugin can be fed back to its input to create repeating sound effects.

The rotation module of this plugin operates in the same way as the [O3A Rotation](#) plugin and can be controlled by an external View application. The reflection, rotation and low pass filter appear *within*

the feedback loop. The wet gain does not, applying only to the final output.

The "Clear Delay" button on the bottom right can be used to clear the audio in the delay line.

12.2.5 View Support

When used with O3A View Sync from the O3A 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 post-rotation 'front' of the scene is shown using a cross. Hold the left mouse button to set all four rotation controls, moving the 'front' of the scene to the indicated direction. Only 2DOF is supported and the rotation will be set so 'up' is upwards.	The scene orientation is shown using a floating cube with matching marker beams on the cube and VR controllers. Rotate a VR controller into the chosen orientation and hold the trigger to set the four rotation controls. Full 3DOF is supported.

12.2.6 Controls

12.2.6.1 Controls: Delay Time and Delay Type

The delay type determines whether the delay time is specified in beats or in seconds. When beats are used, the delay length is a multiple of the beat length in your DAW's current tempo, so if your tempo is in crotchets, then setting the delay to 2.0 results in a delay time equivalent to a minim.

Delay times are rounded to the nearest sample and the minimum delay is one sample. The maximum setting is 20 seconds or 20 beats.

Changes to the delay time are implemented using fades between different delay settings. This means that the delay cannot be used to simulate Doppler shift, but that changes are less likely to be disruptive to musical material. Continuous changes to the delay time may result in artefacts.

Please note that high delay time settings use a significant amount of memory (RAM). This memory is not released until the plugin instance is unloaded.

12.2.6.2 Controls: Dry Gain and Wet Gain

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

12.2.6.3 Control: Feedback Gain

The feedback gain determines how much of the delayed output is sent back into the delay line. The wet gain is not applied here, but rotation and the low pass filter are.

When the feedback gain is set at or close to 0dB, the delay output level is liable to increase over time. Please be careful not to damage anything because of this. The delay line can be cleared using the [Clear Delay](#) button, but being conservative with the feedback gain setting is normally recommended.

12.2.6.4 Control: LPF Cutoff

This controls the cutoff frequency of a simple -6dB/octave low-pass filter applied to the delay output. This can be used to make the sound more distant as echoes repeat. Feedback can result in the filter being applied more than once to the same material.

12.2.6.5 Control: Reflection

This can be used to apply a fixed spatial reflection to the O3A 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.

12.2.6.6 Controls: Pitch, Roll and Yaw

A scene rotation is applied to the delayed O3A scene, in a similar way to the rotation used by the [O3A Rotation](#) plugin. Feedback can result in the rotation being applied more than once to the same material.

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.

12.2.6.7 Control: Ordering

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).

12.2.6.8 Control: Clear Delay

The "Clear Delay" button can be found at the bottom right of the plugin. While pressed, this button clears the delay line and silences everything but the dry signal.

13 Appendix: O3A Streams

13.1 What is an O3A Stream?

A third order ambisonic (O3A) stream is made up of 16 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 16 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 12 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 O3A stream, we find that it is normally best to use an [O3A Visualiser](#) or [O3A Flare](#) plugin.

13.2 What processing can I apply to an O3A stream?

As well as processing designed specifically for O3A 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 16 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 O3A stream with a conventional mono DSP algorithm you *must apply the same processing to all 16 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.

13.3 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.

13.4 How does O3A 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 O3A Core pack includes plugins to do this: [O3A Decoder - FuMa](#) and [O3A Injector - FuMa](#).

13.5 Why Third Order?

These plugins are based on Higher Order Ambisonic (HOA) techniques, which in principle can run at any "order" of spatial detail. We could have made all these plugins work at different orders; once you have got as far as making the maths and DSP work at third order (16 channels) it is not hard to extend to (for instance) fifth order (36 channels). We decided to use third order as it provides a good balance between spatial detail, current practice and CPU load. And sticking to just one type of B-Format avoids all sorts of "wiring" headaches!

That said, if you have a serious need for even higher order algorithms, please get in touch.