O3A Decoding Plugins

v2.4.0

Table of Contents

Table of Contents

1 Introduction

The O3A Decoding plugins decode (or "render") third order ambisonic (O3A) streams for playback on a range of speaker layouts including 5.1, 7.1, 7.1.2 (Dolby Atmos), 22.2, Auro-3D, IMAX and binaural audio for headphones. Other plugins in this collection include converters to other ambisonic formats, a meter, a brickwall limiter and a crosstalk cancellation plugin which complements the headphone decoder.

The speaker and headphone decoders are based on Blue Ripple Sound's ground-breaking Rapture3D Higher Order Ambisonic decoder, running at third order.

This documentation also covers the O3A Core, which is a set of essential tools for work with O3A, and which includes a decoder for basic stereo.

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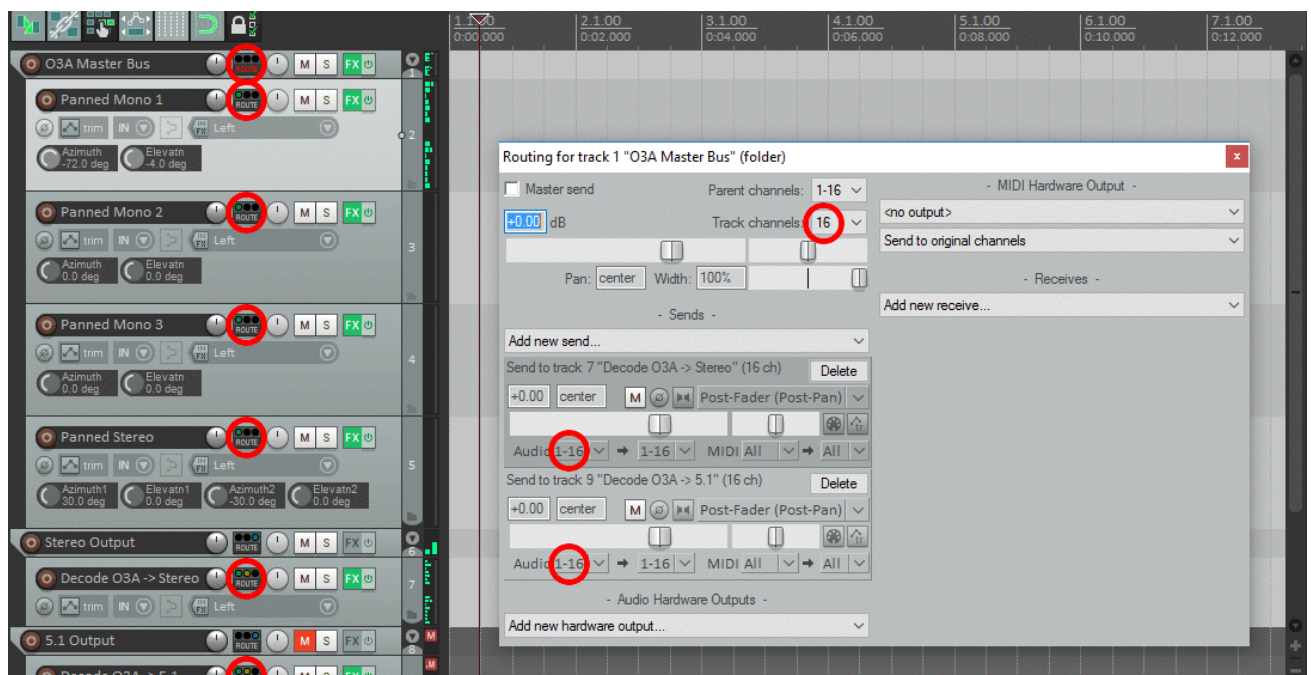
2.2.3 Layout
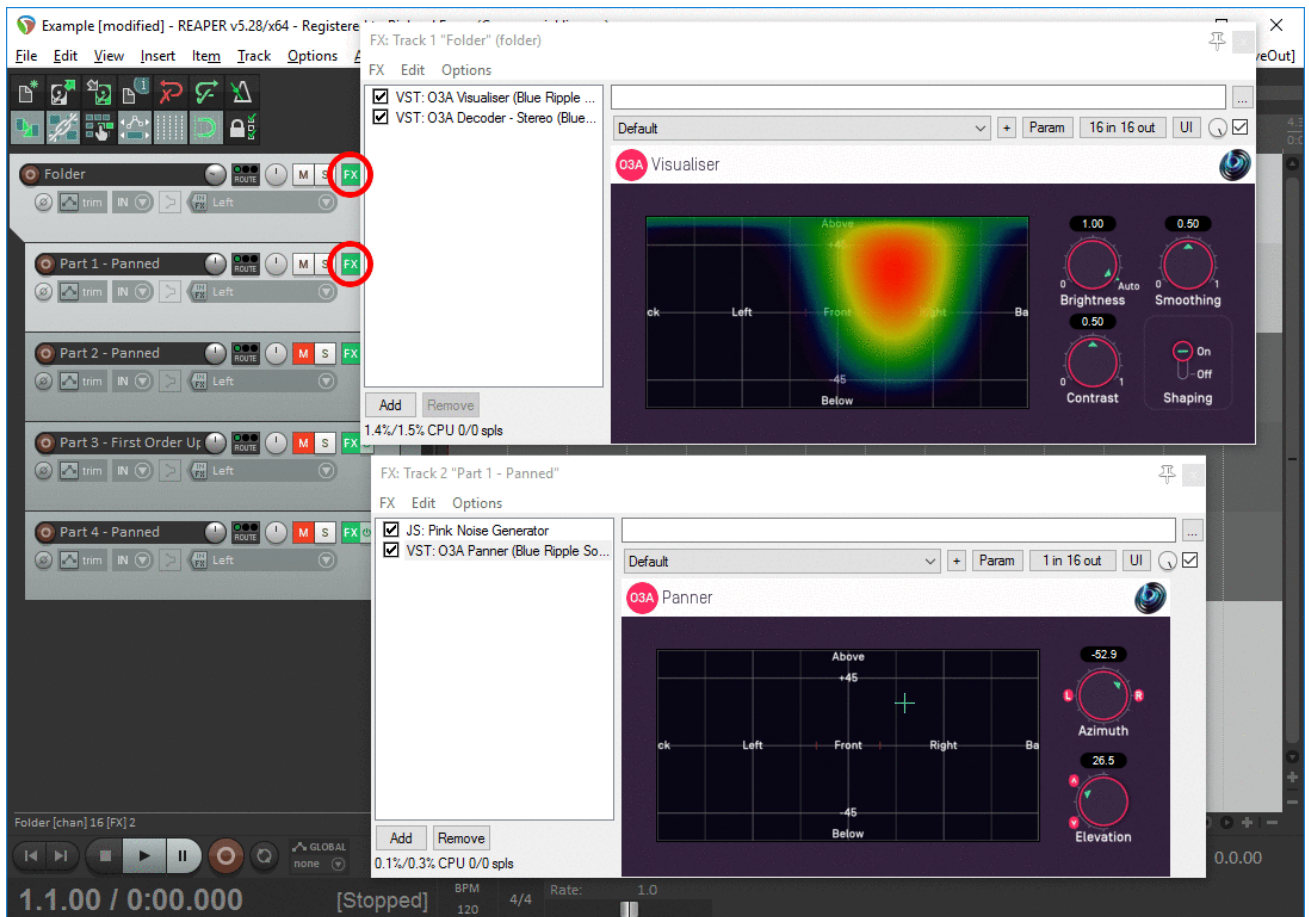
On the folder track, add the following plugins *in this order* to the FX chain:

- 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.
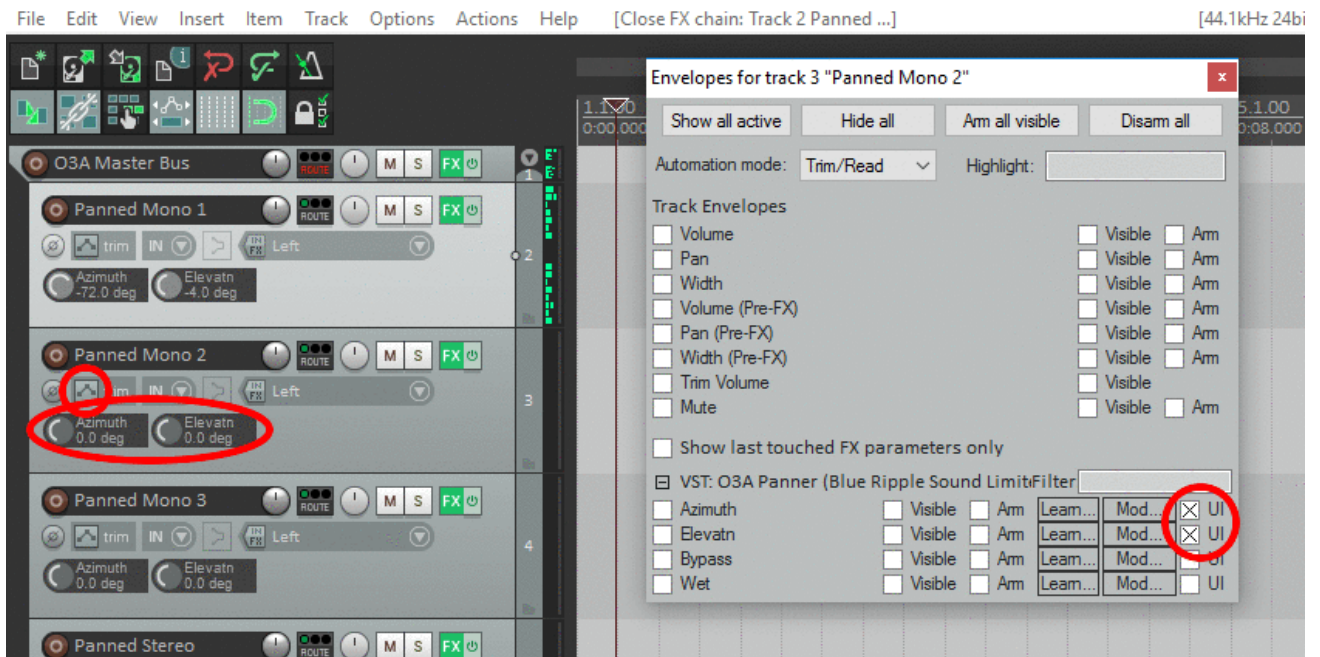
Parts that are already first order ambisonic B-Format can typically be injected directly without issue, although processing can help (e.g. the O3A First Order Injector plugin from the O3A Upmixers set, or the O3A Harpex Upsampler). Of course, other O3A material can be injected directly. FuMa material 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 particularly 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".
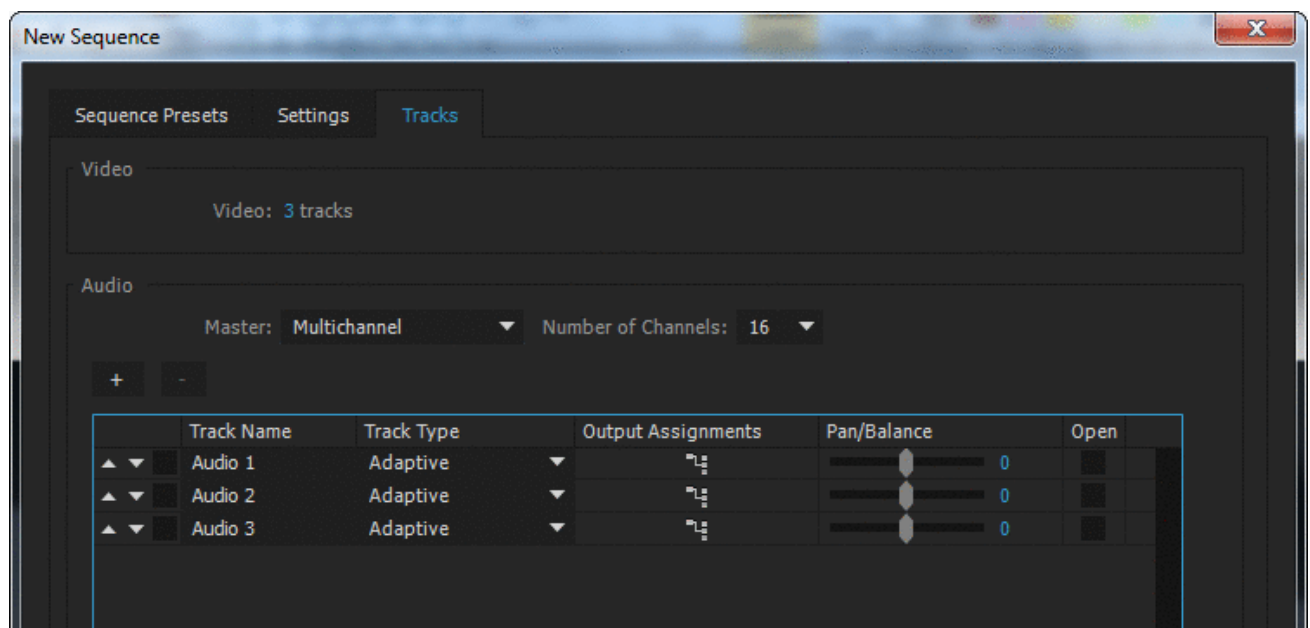
Premiere Pro includes some VR Sequence Presets that support ambisonics. However, these only set the number of channels to 4 by default, which is correct for first order ambisonics, but not for third order. 16 channels are needed by these plugins, so please check the "Number of Channels" as above, or 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.
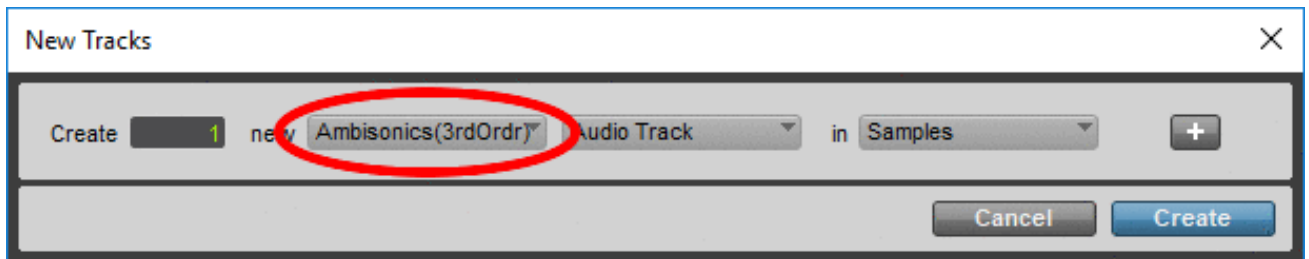
Like these plugins, Pro Tools normally uses the SN3D/ACN ambisonic format. You can convert from FuMa to SN3D using the O3A Injector - FuMa plugin.

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.
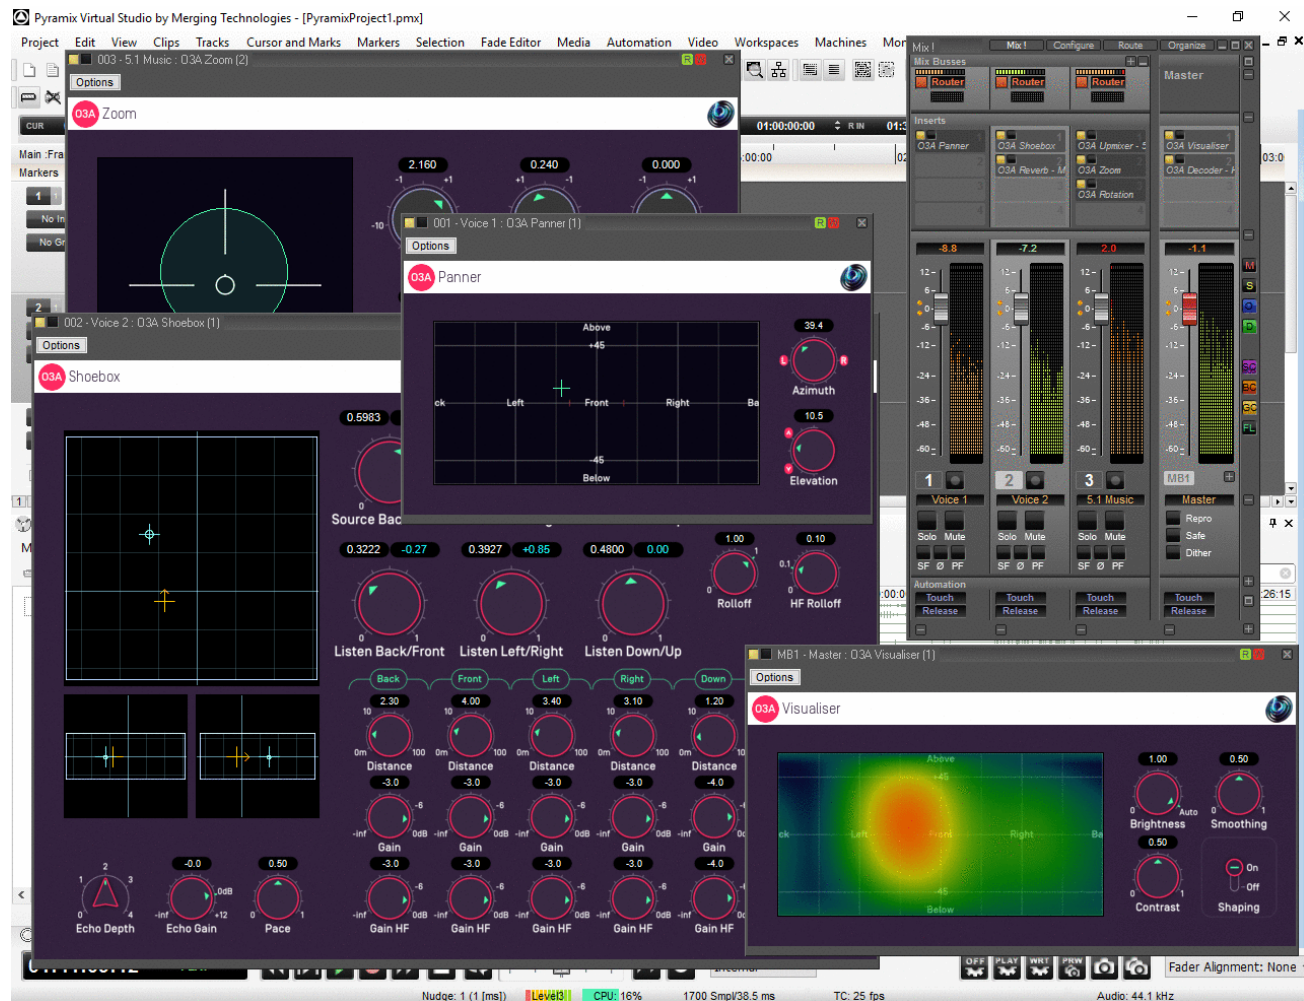
2.5.1 Setup

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).
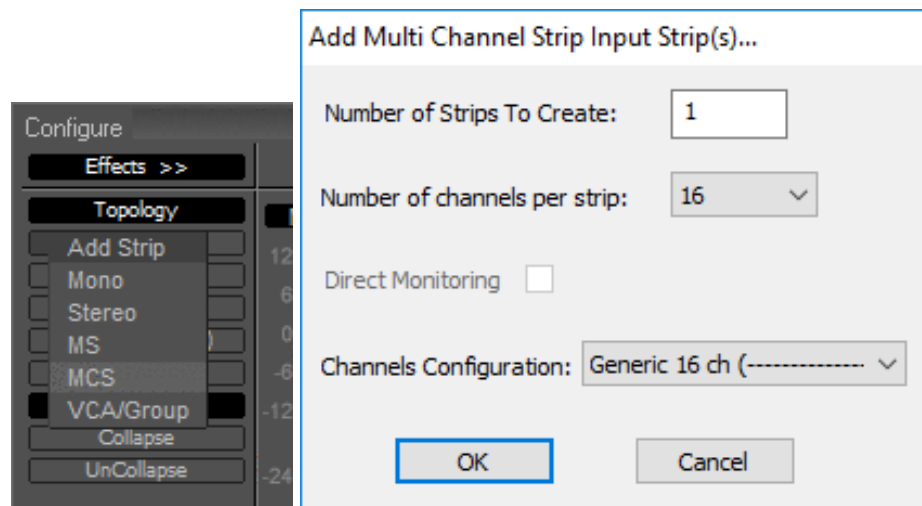
Custom 16 channel mix busses can also be used.

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.
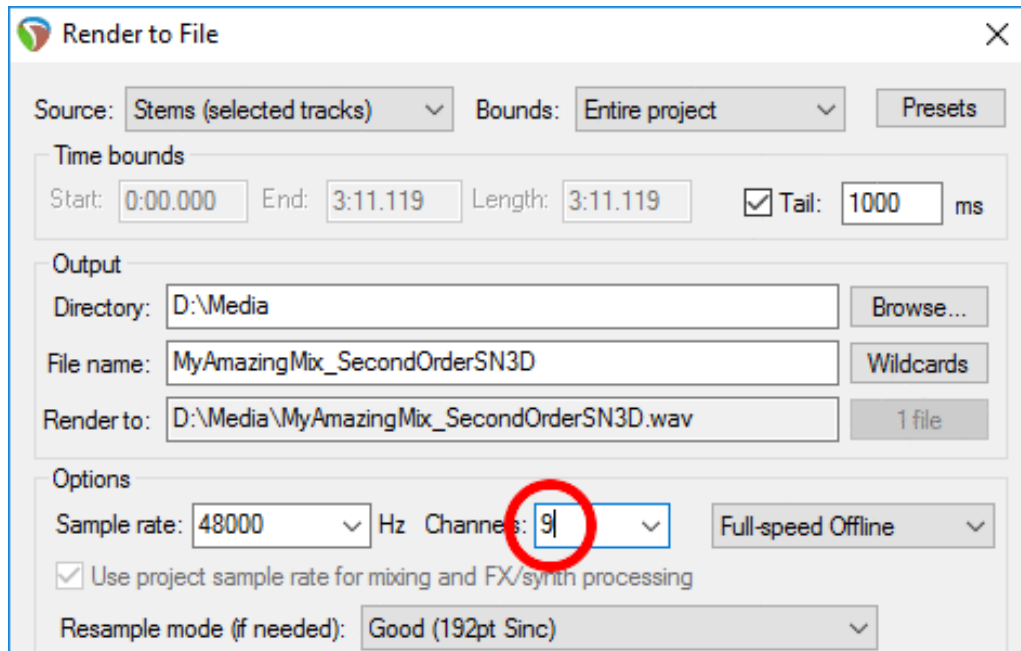
Once you have a nine channel second order file, this can be encoded using the Facebook FB360 Encoder (v1.3.0 or later). Select format "B-format ambiX 2nd order".
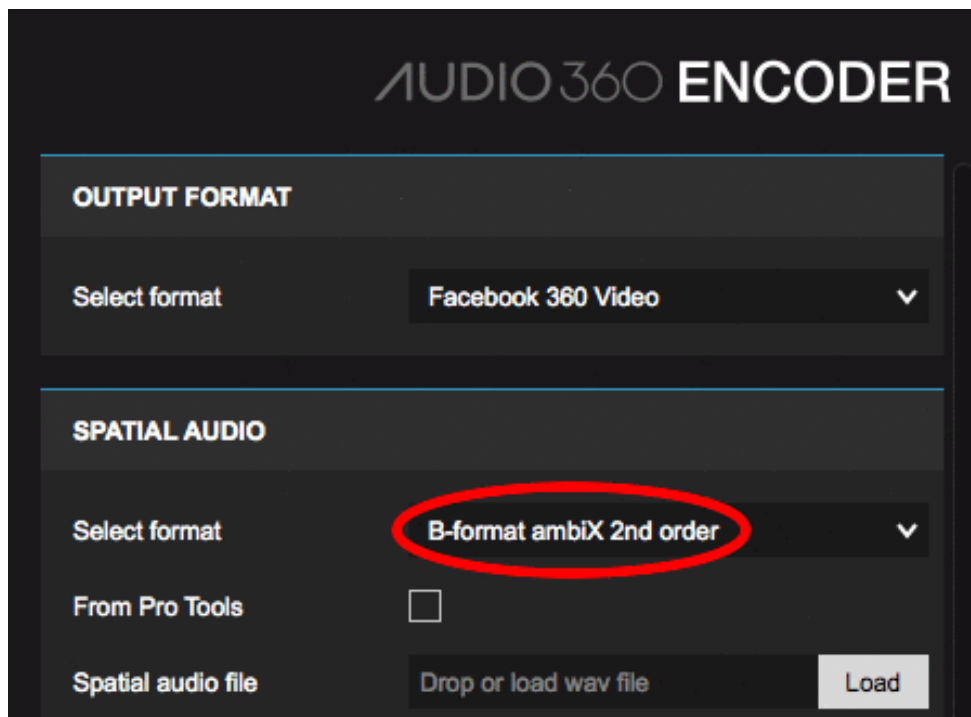
2.6.4 Rapture3D Universal

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.
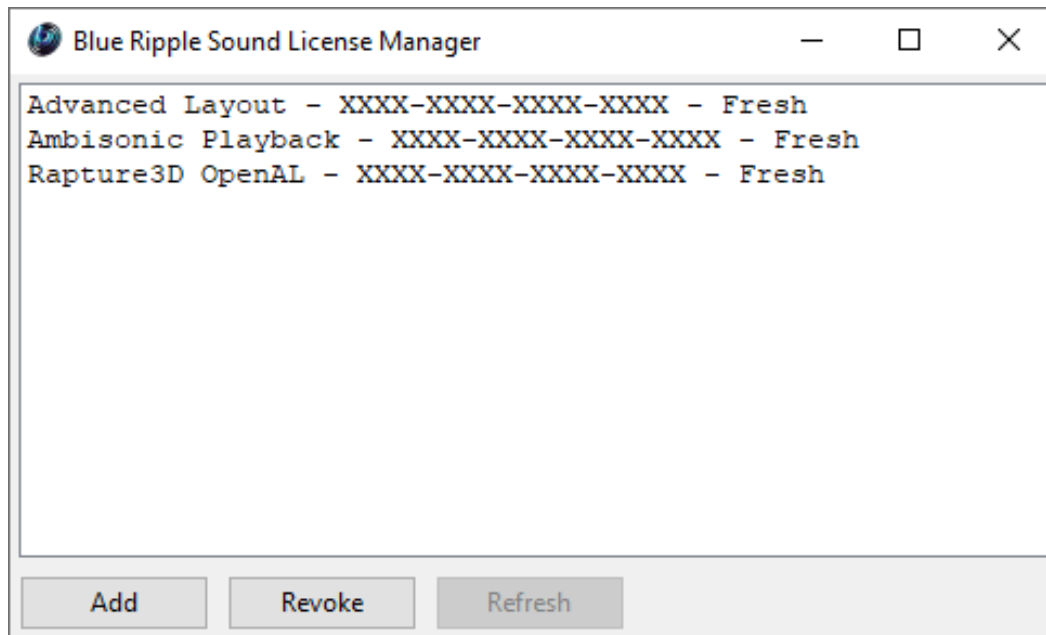
The application is generally available in the Start Menu on Windows, and in your installation directory, which is typically:

- `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.

Due to channel and stem restrictions, not all plugins are supported in AAX. Please see the plugin descriptions for details.

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 Decoding

3.1 When should you decode?

Ideally, decoding should only ever happen in the final listening space. If the O3A stream is sent there it can be decoded with the best algorithm for that particular space and speaker layout. Or, for Virtual Reality, the O3A stream can be rotated to take head orientation into account before decoding, and the decoding might be tailored to the individual's ears, head shape and headphones.

However, it is common that an O3A decoder will not be available in the listening space, in which case decoding needs to happen in the studio. And, in the studio, a decoder of some kind will be needed for monitoring. These plugins can produce stereo, 5.1, 7.1, Auro-3D, IMAX and/or other mixes from your O3A master.

This plugin library provides a number of preset decoders for standard speaker layouts. The "Advanced" edition of Rapture3D from Blue Ripple Sound allows you prepare decoders that are tailored more precisely for your exact speaker locations. It can also build decoders for non-standard, customised speaker layouts, even with large numbers of speakers.

This library also contains other decoders, converters, metering and a limiter that can be useful when exporting O3A streams.

3.2 Choosing a Decoder

There are a number of decoders available here and most support a number of decoding methods.

3.2.1 Headphones

If you are listening on headphones, the binaural decoder O3A Decoder - Headphones will almost certainly work best, although you could also try the basic Stereo decoder (from the O3A Core) for a simpler mix which may play more naturally on ordinary stereo speakers.

3.2.2 Stereo Speakers

The simplest and most robust option here is the basic Stereo decoder from the O3A Core.

An alternative to decoding to stereo in this way is to use a virtual stereo microphone (also from the O3A Core). The angle and response of the virtual microphone capsules can be tuned. Note however that, as with a real stereo microphone, you need to worry about some directions being more prominent than others in the final stereo mix.

If everything is set up very precisely in the listening environment, you might also want to try chaining the headphone decoder with the Stereo Crosstalk Cancellation plugin to produce a wider image. But be aware that this approach has a tiny sweet spot and is not normally suitable for a mix you want to play elsewhere.

3.2.3 Loudspeaker Decoders

Decoders of this type are available for 5.1, 7.1, 7.1.2 (Dolby Atmos), 22.2, Auro-3D, IMAX and others.

Three methods are supported: "Basic", "Reconstruction" and "Tinted Reconstruction".

Generally, "Basic" is recommended for downmixing for distribution and for listening on speaker rigs that are not set up accurately, or are in larger venues. It produces a high quality, resilient "all-purpose" mix.

The other methods can produce a better sound, but are optimised for fairly idealised speaker layouts and listener locations.

If you need a custom layout you may wish to consider the "Advanced" edition of Rapture3D.

3.2.4 Other Ambisonic Formats

These plugins assume a particular method of encoding third order ambisonics generally referred to as "SN3D" (using "ACN" channel ordering). This changed in version 2.0 of these plugins; previously, they used FuMa.

To output in the "FuMa" format, convert using the O3A Decoder - FuMa plugin from the O3A Core.

To output in the "N3D" format (using "ACN" channel ordering), convert using the O3A Decoder - N3D plugin.

To reduce the order, simply use the first four channels (first order) or nine channels (second order) and discard the rest. In Pro Tools, the O3A Decoder - O1A or O2A plugins can simplify this step.

3.3 Low Frequency Effects (LFE)

When decoders in this library manage an output format with a Low Frequency Effect (LFE) channel, a "LFE Crossover" control is shown. The control is off by default.

When this is switched on, some audio is removed from the main channels and routed to the LFE, effectively performing some early bass management. Note that this is in addition to any bass management that may be happening in the playback venue.

We generally do *not* recommend using this. Depending on what happens to the mix, LFE channels may be removed before final playback. Further, audio in this channel is effectively not spatialised. If bass management is needed, it is generally best handled by the playback system, which should also know the crossovers of the actual speakers involved.

However, there are reasons you may need to generate LFE. These include lack of bandwidth, or downstream Quality teams who expect to see some content in this channel! In both these cases, you may find that a very low crossover frequency (such as 10Hz) is satisfactory.

Note that the transition between "On" and "Off" is not necessarily seamless as this switch disables the crossover module to save CPU load.

Audio is split directly, at its original level, so you may need to apply a further gain to the LFE output during some mastering processes (typically -10dB). The "LFE Gain" control can be used for this purpose.

3.4 Decoder Levels and Loudness

All O3A decoders and custom decoders generated using Blue Ripple Sound's Rapture3D "Advanced" edition are designed to produce similar overall output levels.

The O3A Meter or O3A Meter - Karma plugins can be used to estimate these levels from the O3A stream, before decoding. The Karma meter does this in a way that produces numbers which are roughly equivalent to those produced by the ITU-R BS.1770 loudness standard which forms the basis of EBU R128 and ATSC A/85 loudness recommendations and the CALM act.

When exporting an O3A stream to a file, you may wish to use the O3A Brickwall limiter to protect against clipping.

4 O3A Decoding - Binaural Decoding

4.1 O3A Decoder - Headphones
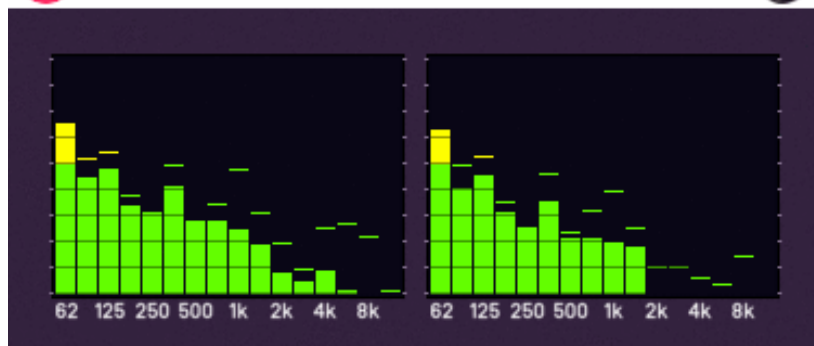
4.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.1.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Binaural Stereo

4.1.3 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it to 3D stereo suitable for use with headphones. The results are similar to a binaural recording and can be used to place sounds behind, around, above and below the listener.

This plugin uses Blue Ripple Sound's ground-breaking "Amber" HRTF technology. This is also available in the Rapture3D game engines and O3A View plugin library.

Real-time plots are shown of the peak frequency content of the output left and right channels, using half-octave frequency bands centred from about 62Hz to 11kHz.

If you are targeting stereo speakers, the output of this plugin may be passed through the Stereo Crosstalk Cancellation plugin, which uses the same HRTF data. Under highly controlled conditions, this can be used with the stereo output from this plugin to produce 3D sound from stereo speakers.

4.1.3.1 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.

The HRTF data for each head in the data set is cleaned and converted into a suitable parameterised model. The parameter space is explored to find an "average" head model which

works on a wide range of heads. This is then used to build a higher order ambisonics (HOA) decoder.

This results in smooth HRTF-based decoders which work on a wide range of heads and do not suffer from the high frequency phasing artefacts common with some other headphone decoders and binaural panners, particularly when sounds are moving. The decoder is able to synthesise ITD (inter-aural time difference) cues at low frequencies and ILD (inter-aural level differences) throughout the frequency range. This plugin uses the third order (O3A) version of the decoder.

5 O3A Decoding - Loudspeaker Decoding

5.1 O3A Decoder - Quad
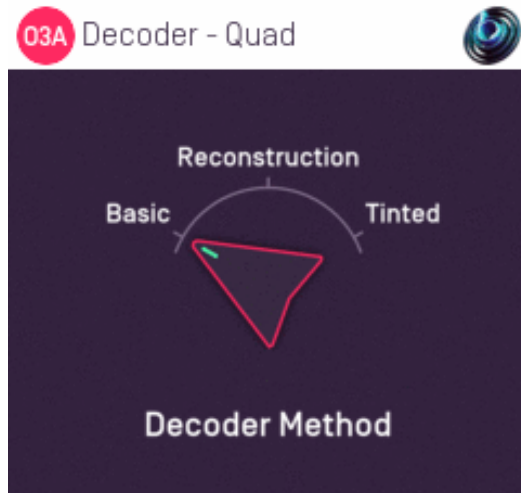
5.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.1.2 Audio

	Channels	Content
Input	16	O3A
Output	4	Quad

5.1.3 Controls

- Method

5.1.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as speaker feeds for four speakers set out at the corners of a square.

5.1.4.1 Channels

The channel ordering used is:

Channel	AAX/VST
1	Front Left
2	Front Right
3	Back Left
4	Back Right

5.1.5 Controls

5.1.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.2 O3A Decoder - 5.1
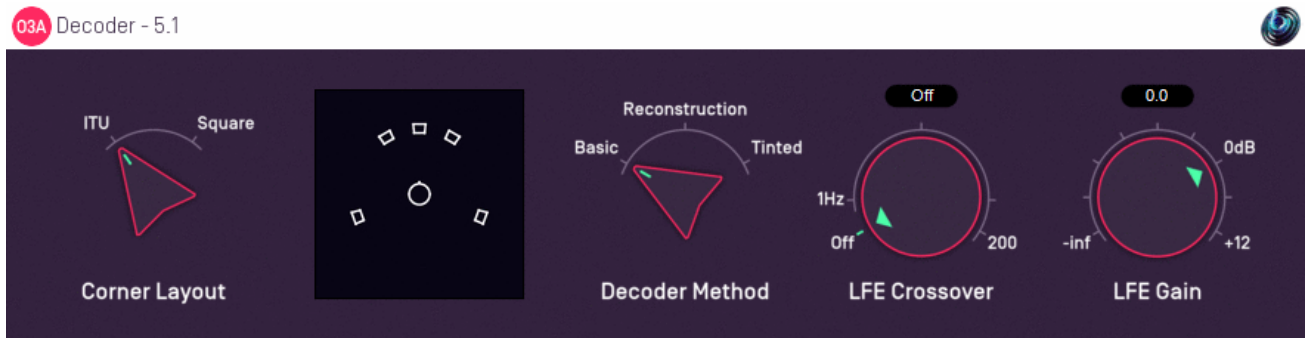
5.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.2.2 Audio

	Channels	Content
Input	16	O3A
Output	6	5.1 Surround

5.2.3 Controls

- Corners
- Method
- LFE Crossover
- LFE Gain

5.2.4 Description

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

5.2.4.1 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

5.2.5 Controls

5.2.5.1 Control: Corners

The corner layout lets you tell the processor where the corner speakers for 5.1 playback are assumed to be:

Corner Layout	Description
ITU	In principle, the speakers of a 5.1 system should be set out using ITU angles. From front centre, this recommends front left and right speakers be at +30 or -30 degrees and side speakers be at +110 or -110 degrees. ITU layouts are common in studios and are a common way to mix 5.1.
Square	In practice, many 5.1 speaker systems are set up with the corner speakers set out in what is roughly a square and some mixes are put together on this basis. If you know that's the case, switch the corner layout to Square.

5.2.5.2 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.2.5.3 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.2.5.4 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.3 O3A Decoder - 7.1

5.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.3.2 Audio

	Channels	Content
Input	16	O3A
Output	8	7.1 Surround

5.3.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.3.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as a 7.1 surround mix.

5.3.4.1 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 Side	LFE
5	Right Surround Side	Back Left
6	Left Surround Rear	Back Right
7	Right Surround Rear	Side Left
8	LFE	Side Right

5.3.5 Controls

5.3.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.3.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.3.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.4 O3A Decoder - 7.1.2 (Dolby Atmos)

5.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.4.2 Audio

	Channels	Content
Input	16	O3A
Output	10	7.1.2 (Dolby Atmos)

5.4.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.4.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as 7.1.2.

7.1.2 is a standard mix format for beds in Dolby Atmos. It adds two ceiling channels to standard 7.1.

5.4.4.1 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 Side	LFE
5	Right Surround Side	Left Surround Rear
6	Left Surround Rear	Right Surround Rear
7	Right Surround Rear	Left Surround Side
8	LFE	Right Surround Side
9	Left Top Surround	Left Top Surround
10	Right Top Surround	Right Top Surround

Please note that the channel ordering used by Dolby's Atmos RMU may differ. Also, the AAX order above is used when plugins interact with ProTools through the AAX software interface; other parts of ProTools may use a different default order, for instance when exporting to file.

5.4.5 Controls

5.4.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

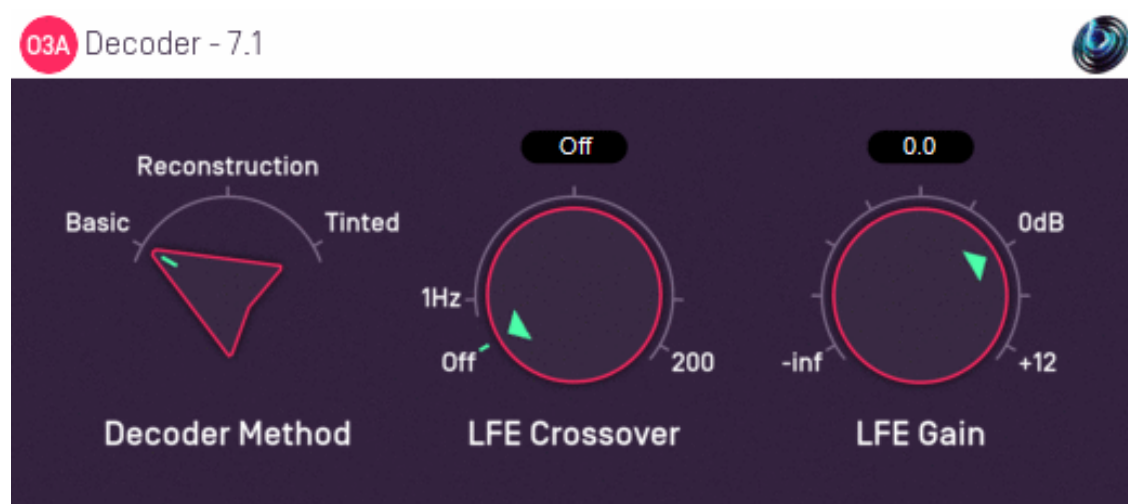| Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.4.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.4.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.5 O3A Decoder - 7.1.4

5.5.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.5.2 Audio

	Channels	Content
Input	16	O3A
Output	12	7.1.4 Surround

5.5.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.5.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it to 7.1.4.

The AAX version of this specific plugin requires Pro Tools 2023.6 or later.

Please note that although 7.1.4 is not an uncommon speaker layout for Dolby Atmos playback, the actual mastering and delivery format used is often 7.1.2.

5.5.4.1 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 Side	LFE
5	Right Surround Side	Left Surround Rear
6	Left Surround Rear	Right Surround Rear
7	Right Surround Rear	Left Surround Side
8	LFE	Right Surround Side
9	Left Top Front	Left Top Front
10	Right Top Front	Right Top Front
11	Left Top Rear	Left Top Rear
12	Right Top Rear	Right Top Rear

5.5.5 Controls

5.5.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.5.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.5.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.6 O3A Decoder - Auro-3D 9.1, 10.1, 11.1 and 13.1

5.6.1 Host Support

5.6.1.1 O3A Decoder - Auro-3D 9.1

Host Type	Support
AAX	Yes, output mapped to 5.1.4
VST2	Yes

5.6.1.2 O3A Decoder - Auro-3D 10.1

Host Type	Support
AAX	Yes, output mapped to 7.0.4
VST2	Yes

5.6.1.3 O3A Decoder - Auro-3D 11.1

Host Type	Support
AAX	Yes, output mapped to 7.1.4
VST2	Yes

5.6.1.4 O3A Decoder - Auro-3D 13.1

Host Type	Support
AAX	Yes, output mapped to 7.1.6
VST2	Yes

5.6.2 Audio

5.6.2.1 O3A Decoder - Auro-3D 9.1

	Channels	Content
Input	16	O3A
Output	10	Auro-3D 9.1 (mapped to 5.1.4 for AAX)

5.6.2.2 O3A Decoder - Auro-3D 10.1

	Channels	Content
Input	16	O3A
Output	11	Auro-3D 10.1 (mapped to 7.0.4 for AAX)

5.6.2.3 O3A Decoder - Auro-3D 11.1

	Channels	Content
Input	16	O3A
Output	12	Auro-3D 11.1 (mapped to 7.1.4 for AAX)

5.6.2.4 O3A Decoder - Auro-3D 13.1

	Channels	Content
Input	16	O3A
Output	14	Auro-3D 13.1 (mapped to 7.1.6)

5.6.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.6.4 Description

Four decoder plugins are provided, which produce Auro-3D surround mixes from O3A streams. The plugins correspond to each of the Auro-3D 9.1, Auro-3D 10.1, Auro-3D 11.1 and Auro-3D 13.1 speaker layouts.

The AAX versions of these specific plugins require Pro Tools 2023.6 or later.

These plugins do *not* perform bitstream encoding to allow transmission of Auro-3D over backwards-compatible formats. The output is in a form suitable to be fed directly to speaker channels, or to such an encoder.

5.6.4.1 Channels

The channel orderings used for 9.1 use the following table. AAX is mapped to the 5.1.4 speaker layout.

Channel	AAX	VST
1	Front Left	Front Left
2	Front Centre	Front Right
3	Front Right	Front Centre
4	Surround Left	LFE
5	Surround Right	Surround Left
6	LFE	Surround Right
7	Height Front Left	Height Front Left
8	Height Front Right	Height Front Right
9	Height Rear Left	Height Rear Left
10	Height Rear Right	Height Rear Right

The other Auro-3D formats are:

Auro-3D 10.1	Auro-3D 11.1	Auro-3D 13.1	Channel Name
1	1	1	Front Left
2	2	2	Front Right
3	3	3	Front Centre
4	4	4	LFE
5	5	5	Surround Left
6	6	6	Surround Right
-	-	7	Back Left
-	-	8	Back Right
7	7	9	Height Front Left
8	8	10	Height Front Right
9	9	11	Height Surround Left
10	10	12	Height Surround Right
-	11	13	Height Front Centre
11	12	14	Top Ceiling

When these formats are mapped for AAX, no attempt is made to map to channels where they correspond. The channels simply appear in the order above. The AAX stem formats actually used are 7.0.4, 7.1.4 and 7.1.6.

5.6.5 Controls

5.6.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.6.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.6.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.7 O3A Decoder - IMAX 5.0, 6.0 and 12.0

5.7.1 Host Support

5.7.1.1 O3A Decoder - IMAX 5.0

Host Type	Support
AAX	Yes, output mapped to standard 5.0
VST2	Yes

5.7.1.2 O3A Decoder - IMAX 6.0

Host Type	Support
AAX	Yes, output mapped to standard 6.0
VST2	Yes

5.7.1.3 O3A Decoder - IMAX 12.0

Host Type	Support
AAX	Yes, output mapped to 7.1.4
VST2	Yes

5.7.2 Audio

5.7.2.1 O3A Decoder - IMAX 5.0

	Channels	Content
Input	16	O3A
Output	5	IMAX 5.0 (mapped to standard 5.0 for AAX)

5.7.2.2 O3A Decoder - IMAX 6.0

	Channels	Content
Input	16	O3A
Output	6	IMAX 6.0 (mapped to standard 6.0 for AAX)

5.7.2.3 O3A Decoder - IMAX 12.0

	Channels	Content
Input	16	O3A
Output	12	IMAX 12.0 (mapped to 7.1.4 for AAX)

5.7.3 Controls

- Method

5.7.4 Description

These plugins take a third order ambisonics (O3A) mix and decode it as IMAX surround mixes.

The AAX version of the IMAX 12.0 plugin requires Pro Tools 2023.6 or later.

5.7.4.1 IMAX 5.0 Channels

The channel ordering used for IMAX 5.0 is as follows:

Channel	AAX	VST
1	Front Left	Front Left
2	Front Centre	Front Right
3	Front Right	Front Centre
4	Rear Left (mapped to Left Surround)	Rear Left
5	Rear Right (mapped to Right Surround)	Rear Right

For AAX, note that this is mapped to standard 5.0.

5.7.4.2 IMAX 6.0 Channels

Channel	AAX	VST
1	Front Left	Front Left
2	Front Centre	Front Right
3	Front Right	Front Centre
4	Rear Left (mapped to Left Surround)	Front Centre Overhead
5	**Front Centre Overhead (mapped to Centre Surround)**	Rear Left
6	Rear Right (mapped to Right Surround)	Rear Right

For AAX, note that this is mapped to standard 6.0, with the IMAX Front Centre Overhead channel mapped to the Centre Surround in standard 6.0 (i.e. at the back).

5.7.4.3 IMAX 12.0 Channels

Channel	VST
1	Front Left
2	Front Right
3	Front Centre
4	Front Centre Overhead
5	Rear Left
6	Rear Right
7	Side Left
8	Side Right
9	Overhead Front Left
10	Overhead Front Right
11	Overhead Rear Left
12	Overhead Rear Right

IMAX systems use speakers with wider bandwidth than typical cinema systems. No Low Frequency Effect (LFE) channel is used in the mix formats. Also note that the last two channels in IMAX 5.0 and 6.0 correspond to rear speakers, not side or surround speakers.

5.7.5 Controls

5.7.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.8 O3A Decoder - 22.2

5.8.1 Host Support

Host Type	Support
AAX	No
VST2	Yes

5.8.2 Audio

	Channels	Content
Input	16	O3A
Output	24	Hamasaki 22.2

5.8.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.8.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as a 24 channel Hamasaki 22.2 surround mix.

5.8.4.1 Channels

The channel ordering used is:

Channel	VST
1	Front Left
2	Front Right
3	Front Centre
4	LFE Left
5	Back Left
6	Back Right
7	Front Left/Centre
8	Front Right/Centre
9	Back Centre
10	LFE Right
11	Side Left
12	Side Right
13	Top Front Left
14	Top Front Right
15	Top Front Centre
16	Top Centre
17	Top Back Left
18	Top Back Right
19	Top Side Left
20	Top Side Right
21	Top Back Centre
22	Bottom Front Centre
23	Bottom Front Left
24	Bottom Front Right

Note that 22.2 needs 24 channels in total, which is even more than the 16 that O3A normally needs. *Don't forget to check your DAW channel count.*

5.8.5 Controls

5.8.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.8.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.8.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

Please note that a further -6dB is applied in addition to this setting to compensate for the presence of two LFEs in the 22.2 configuration.

5.9 O3A Decoder - 3D7.1

5.9.1 Host Support

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

5.9.2 Audio

	Channels	Content
Input	16	O3A
Output	8	3D7.1 Surround (mapped to 7.1 for AAX)

5.9.3 Controls

- Method
- LFE Crossover
- LFE Gain

5.9.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as a 3D7.1 surround mix.

3D7.1 is a modified version of the 7.1 speaker layout invented by Simon N. Goodwin. It is based on the equipment used for 7.1 but places the speakers in different locations, some being raised or lowered.

If you intend to use this decoder then you should move your speakers to these new locations. Alternatively there is a normal O3A decoder for 7.1.

5.9.4.1 Channel Ordering for AAX

For AAX, the channel ordering and assumed speaker locations for the 3D7.1 surround mix are as follow.

Channel	7.1 Speaker Position	3D7.1 Speaker Position	Azimuth	Elevation	Used in 5.1
1	Front Left	Front Left High	+51	+24	Yes
2	Front Centre	Front Centre	0	0	Yes
3	Front Right	Front Right High	-51	+24	Yes
4	Left Surround Side	Rear Left Low	+129	-24	Yes
5	Right Surround Side	Rear Right Low	-129	-24	Yes
6	Left Surround Rear	Rear Centre High	+180	+55	No
7	Right Surround Rear	Front Centre Low	0	-55	No
8	LFE	LFE	-	-	Yes

5.9.4.2 VST

For VST, the channel ordering and assumed speaker locations for the 3D7.1 surround mix are as follow.

Channel	7.1 Speaker Position	3D7.1 Speaker Position	Azimuth	Elevation	Used in 5.1
1	Front Left	Front Left High	+51	+24	Yes
2	Front Right	Front Right High	-51	+24	Yes
3	Front Centre	Front Centre	0	0	Yes
4	LFE	LFE	-	-	Yes
5	Rear Left	Rear Centre High	+180	+55	No
6	Rear Right	Front Centre Low	0	-55	No
7	Side Left	Rear Left Low	+129	-24	Yes
8	Side Right	Rear Right Low	-129	-24	Yes

In the tables above, azimuth is measured anticlockwise (left) from the front and elevation upwards from the horizontal.

5.9.5 Controls

5.9.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.9.5.2 Control: LFE Crossover

This determines if any audio is sent to the LFE channel. It may be set to "Off" and this is recommended. If enabled, audio is extracted from the main mix into the LFE channel using the indicated crossover frequency (between 1 and 200Hz). You may wish to read our notes on this for more detail.

5.9.5.3 Control: LFE Gain

If the LFE gain is active, this gain is applied to it. Some export formats may need this set to -10dB to take into account a +10dB gain applied during playback.

5.10 O3A Decoder - Cube
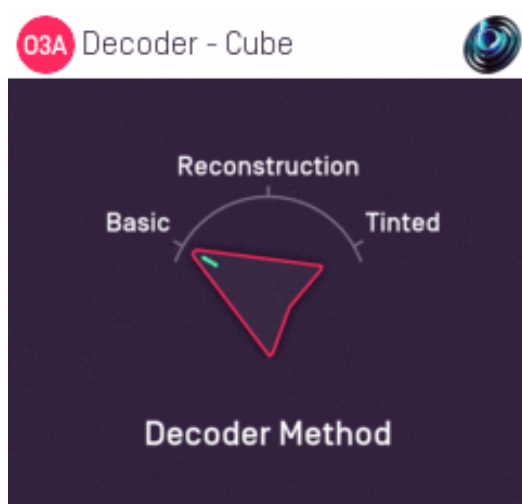
5.10.1 Host Support

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

5.10.2 Audio

	Channels	Content
Input	16	O3A
Output	8	Speaker feeds for cube (mapped to 7.1 for AAX)

5.10.3 Controls

- Method

5.10.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as speaker feeds for eight speakers set out at the corners of an cube.

5.10.4.1 Channels

The channel ordering used is:

Channel	AAX/VST
1	Lower Front Left
2	Lower Front Right
3	Lower Back Right
4	Lower Back Left
5	Upper Front Left
6	Upper Front Right
7	Upper Back Right
8	Upper Back Left

Note that these are mapped directly onto a 7.1 output in the AAX version of the plugin.

5.10.5 Controls

5.10.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.11 O3A Decoder - Hexagon
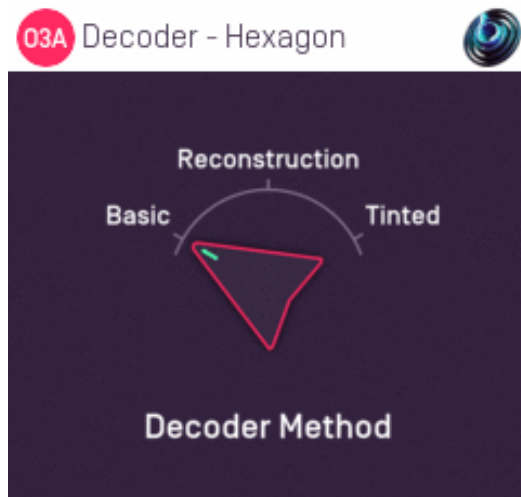
5.11.1 Host Support

Host Type	Support
AAX	Yes, output mapped to 6.0
VST2	Yes

5.11.2 Audio

	Channels	Content
Input	16	O3A
Output	6	Speaker feeds for hexagon (mapped to 6.0 for AAX)

5.11.3 Controls

- Method

5.11.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as speaker feeds for six speakers set out at the corners of an hexagon.

5.11.4.1 Channels

The channel ordering for the speaker feeds assumes that the speakers are set out clockwise around the room, starting at the front and slightly to the left, so the first two channels are roughly where normal "Front Left" and "Front Right" speakers would be.

The channel ordering and speaker azimuths are:

Channel	AAX/VST	Azimuth
1	Front Left	+30
2	Front Right	-30
3	Side Right	-90
4	Back Right	-150
5	Back Left	+150
6	Side Left	+90

In the table above, azimuth is measured anticlockwise (left) from the front. Note that these are mapped directly onto a 6.0 output in the AAX version of the plugin.

5.11.5 Controls

5.11.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

5.12 O3A Decoder - Octagon
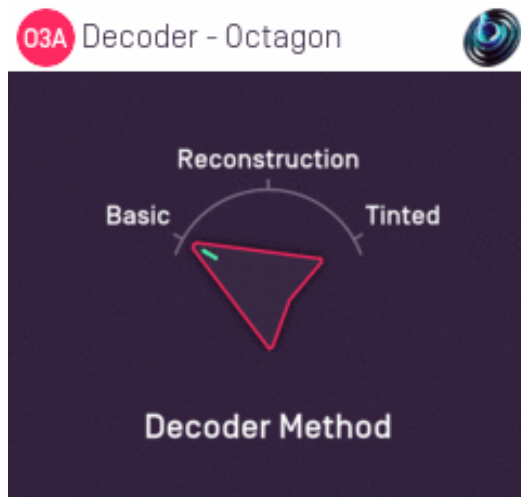
5.12.1 Host Support

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

5.12.2 Audio

	Channels	Content
Input	16	O3A
Output	8	Speaker feeds for octagon (mapped to 7.1 for AAX)

5.12.3 Controls

- Method

5.12.4 Description

This plugin takes a third order ambisonics (O3A) mix and decodes it as speaker feeds for eight speakers set out at the corners of an octagon.

5.12.4.1 Channels

The channel ordering for the speaker feeds assumes that the speakers are set out clockwise around the room, starting at the front and slightly to the left, so the first two channels are roughly where normal "Front Left" and "Front Right" speakers would be.

The channel ordering and speaker azimuths are:

Channel	AAX/VST	Azimuth
1	Left at Front	+22.5
2	Right at Front	-22.5
3	Front at Right	-67.5
4	Back at Right	-112.5
5	Right at Back	-157.5
6	Left at Back	+157.5
7	Back at Left	+112.5
8	Front at Left	+67.5

In the table above, azimuth is measured anticlockwise (left) from the front. Note that these are mapped directly onto a 7.1 output in the AAX version of the plugin.

5.12.5 Controls

5.12.5.1 Control: Method

Three methods are supported, "Basic", "Reconstruction" and "Tinted Reconstruction". Generally, "Basic" is recommended except where the speaker rig to be used for playback is known and has been set up accurately.

Method	Description
Basic	This method produces a resilient decoding which should play well on a range of speaker systems including large venues.
Reconstruction	This method controls the soundfield at the centre of the listening space to produce a clearer spatial image. It relies on the speaker layout being set up accurately.
Tinted Reconstruction	This variant of the Reconstruction method performs HRTF-based tinting on the audio stream to strengthen the spatial image. Again, it relies on the speaker layout being set up accurately.

6 O3A Decoding - Other Decoders and Converters

6.1 O3A Decoder - Quad Binaural
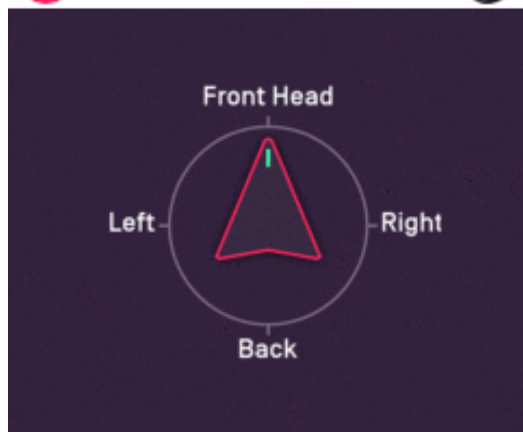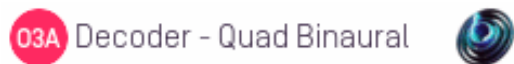
6.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.1.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Binaural Stereo (one head of four)

6.1.3 Controls

- Head

6.1.4 Description

This plugin takes a third order ambisonic (O3A) mix and decodes it to synthetic "Quad Binaural".

Quad Binaural is a format used in some Virtual Reality (VR) applications. It uses four binaural stereo recordings for pairs of ears at ninety degree angles (so eight audio channels in total). The stereo recordings are mixed together during playback in a way that can give the impression of a horizontal scene rotation, typically in response to head-tracking.

We do not particularly recommend the Quad Binaural format for VR applications. If your VR audio engine supports ambisonics properly then better rotational results can be achieved using ambisonic formats directly (for example, in Rapture3D Universal).

The decoder plugin assumes that at any one time during playback, two of the four binaural stereo recordings will be mixed together to produce the final binaural output, using sine and cosine gains (i.e. an equal-power crossfade). This can cause various artefacts which the binaural processing used in this plugin attempts to compensate for. Because of this, when only one head is being used,

the results are typically *not* as good as when using the standard headphone decoder directly.

We recommend use of the standard binaural headphone decoder (O3A Decoder - Headphones) instead of this plugin when Quad Binaural is not in use.

6.1.5 Controls

6.1.5.1 Control: Head

This control identifies which of the four binaural stereo recordings should be synthesized. They are:

Head	Scene Rotation
Front	0 degrees
Right	90 degrees
Back	180 degrees
Left	270 degrees

For instance, the "Right" head corresponds to the case where the head is looking 90 degrees to the right. The scene is thus rotated 90 degrees to the left, so a sound that was previously at the front will be heard on the left.

6.1.5.2 Technical Notes

This plugin is based on the Blue Ripple Sound Amber HRTF (see the O3A Decoder - Headphones plugin). However, it has been simplified and modified to improve results with Quad Binaural rotation.

The plugin compensates assuming a sine and cosine equal-power crossfade during rotation. Results with linear crossfades are not expected to be as good.

6.2 O3A Decoder - UHJ Stereo

6.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.2.2 Audio

	Channels	Content
Input	16	O3A
Output	2	UHJ Stereo

6.2.3 Description

This plugin takes a third order ambisonics (O3A) mix and converts it into UHJ stereo. This form of stereo uses phase to encode a small amount of ambisonic surround sound information, which can be recovered later, for instance in special A/V receivers.

UHJ stereo was introduced in 1985 for use with first order ambisonics. It contains no height information, only takes audio from three of the sixteen channels that make up O3A and those channels cannot be recovered entirely. However, it does allow an impressive amount of spatial detail to be "hidden" in the stream while still being good to listen to in stereo.

UHJ stereo is sometimes used for artistic reasons, as stereo formed this way has particular sound qualities. For instance, UHJ produces a "phasey" stereo image when sounds are behind, which can be desirable.

For simple stereo decoding you might also wish to try the basic O3A Decoder - Stereo plugin from the O3A Core plugin library. The O3A Virtual Microphone from the same library also produces stereo.

This plugin uses a gain boost of +4.315dB relative to the common UHJ decoding equations, for consistency with other Blue Ripple Sound studio decoders.

6.3 O3A Decoder - N3D
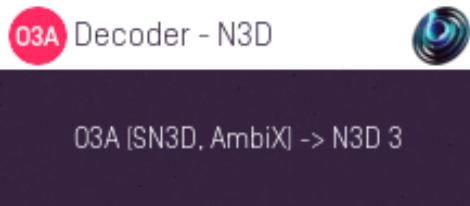
6.3.1 Host Support

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

6.3.2 Audio

	Channels	Content
Input	16	O3A
Output	16	N3D/ACN Ambisonic Audio (mapped to standard SN3D ambisonics for AAX)

6.3.3 Description

This plugin takes a third order ambisonics (O3A) mix and converts it to a third order "N3D" ambisonic mix. N3D is a close relative to SN3D, which is used natively by these plugins. The "ACN" channel ordering convention is used in both cases.

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.

In Pro Tools, as the N3D format is not supported directly, 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.

N3D is used in MPEG-H 3D Audio.

6.3.3.1 Warning

Please be aware that N3D and SN3D are similar to each other, but not the same. Be **very** careful to convert where necessary as it may not be immediately obvious if mistakes are made. **Do *not* guess!** If you get this wrong it is likely to result in level, blurring or sharpening errors which may not be immediately obvious on simple decoders.

6.3.3.2 Other Conventions

Note that other plugins are available to convert to and from ambisonics using the FuMa convention. See O3A Decoder - FuMa and O3A Injector - FuMa.

6.3.3.3 Orders

This plugin produces third order N3D 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 from the O3A Core can be "misused" to reduce the order of the N3D material as this is a rare case where a plugin will work correctly for both SN3D and N3D.

7 O3A Decoding - Utilities

7.1 O3A Brickwall
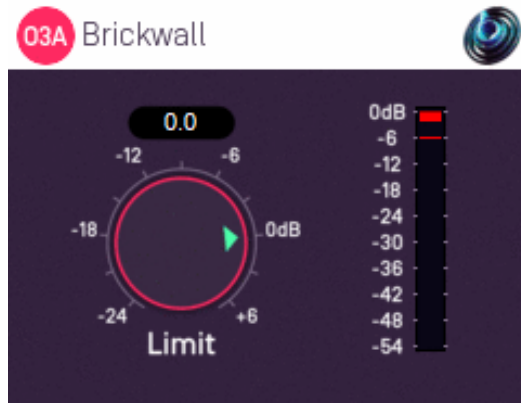
7.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

7.1.3 Controls

- Limit

7.1.4 Description

This plugin ensures that the amplitude of all channels in the signal are below a particular limit. It does this by applying a "brick wall" limiter with a fast attack to the signal. The limiter is triggered when the signal is within a decibel of the target level (or above). When the limiter is triggered it reduces the level of the signal aggressively, reducing all channels by the same gain. When the limiter is not triggered the signal passes through unmodified.

A meter shows how much the level is currently reduced, along with the extreme used in the last ten seconds.

For instance, this can be used for safety when exporting to a fixed-point PCM file, to ensure that no signals clip. The sound may have artefacts but generally these are far better than clipping.

If you are looking for dynamic range control plugins for artistic reasons, you may also want to consider the O3A Compressor from the O3A Manipulators plugin library.
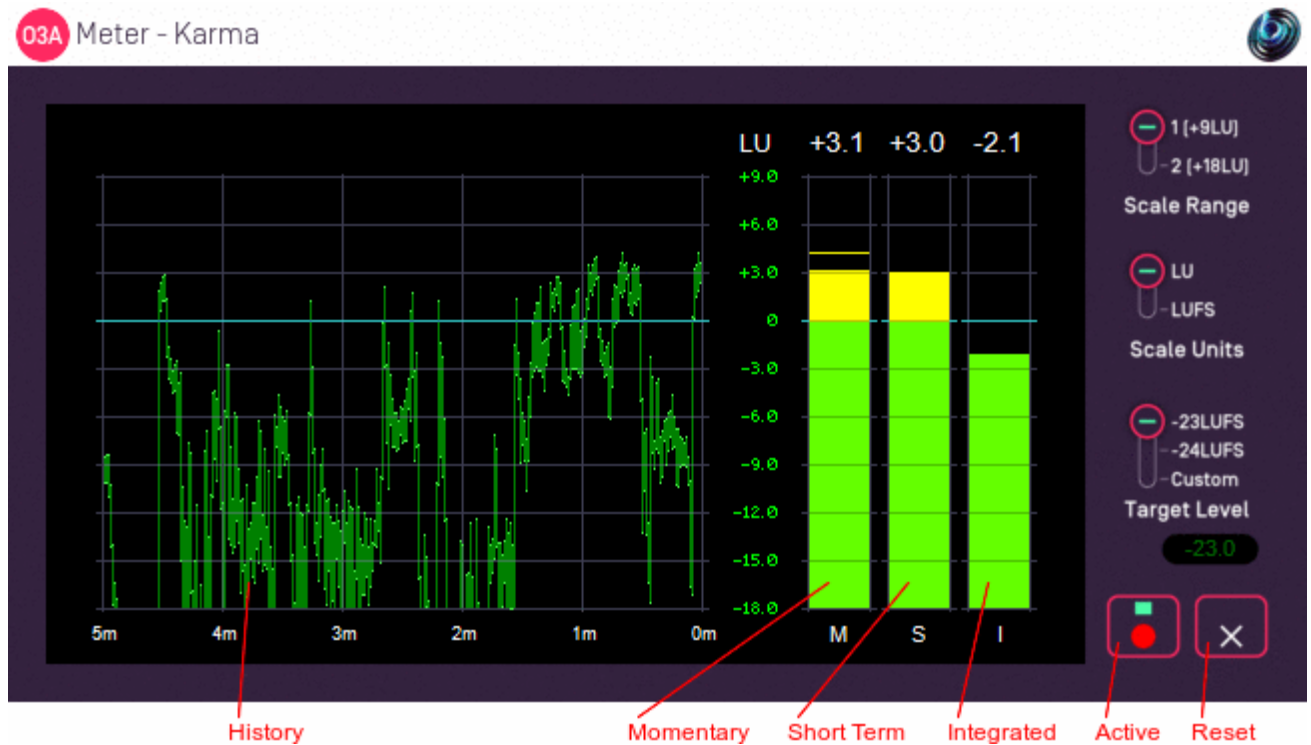
7.1.5 Controls

7.1.5.1 Control: Limit

This control determines the level above which the output may not go. Values are in decibels and a value of 0dB corresponds to a "full scale" signal and is the default. Please note that this convention is *not* the convention used by the O3A metering plugins, so values will not correspond directly.

Note that subsequent external processing such as sample rate conversion can result in higher signal peaks. To allow a little extra headroom, you may wish to reduce the limit value by a few decibels.

7.2 O3A Meter - Karma

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

- Scale Range
- Scale Units
- Target Level
- Active
- Reset

7.2.4 Description

This plugin measures the loudness of an O3A stream. It provides metering and passes audio through unchanged.

Its design is intended to be a generalisation of the ITU-R BS.1770 standard loudness algorithm which forms the basis of the EBU R128 and ATSC A/85 loudness recommendations and the CALM act, but applied to an O3A stream rather than a stream for a specific speaker layout like 5.1.

The O3A decoders included in the O3A Core and O3A Decoding libraries and custom decoders generated using Blue Ripple Sound's Rapture3D "Advanced" edition are all designed in a way that means that the perceived loudness after decoding will be reasonably consistent with the numbers produced by this meter.

This means that if you use this plugin to meter your O3A mix, it will give you a good idea of the loudness you will achieve later on when you decode for a particular speaker layout.

7.2.4.1 Modifying ITU-R BS.1770 for O3A

The algorithm used for metering here is based on the ITU-R BS.1770 loudness algorithm. It uses the same "K" frequency weighting, RMS calculation, windows and gating, although a low pass filter is applied to the K-weighting above 24kHz when the sample rate is above 48kHz.

However, BS.1770 is specified for a 5.1 mix only whereas this meter is designed for use with full 3D O3A streams. This distinction means that channel contributions are aggregated using a substantially different methodology. Specifically, a smooth front/back directional weighting is used in place of the +1.5dB applied to the 5.1 surround channels in the standard. This weighting is approximately 0dB at the front and +1.5dB at +/-110deg (the "surround" angles of 5.1).

This distinction means that results are not identical to those produced by decoding to headphones, stereo, quad and 5.1 and then applying a standard BS.1770 meter. However, if Blue Ripple Sound O3A or custom Rapture3D decoders are to be used, this meter will provide a good estimate without the need to run each decoder and check individual results.

This meter does not perform True Peak measurements. Note that True Peaks can vary hugely depending on the decoder chosen.

7.2.4.2 "Different" Loudness Standards

EBU R128 and ATSC A/85 are both based on the ITU-R BS.1770 loudness algorithm but there are some subtle differences.

The full-scale unit for loudness is called "LKFS" in BS.1770 and A/85, but "LUFS" in R128. However, these are actually two names for the same thing. We use LUFS. A 0dBFS 1kHz sine tone panned to the front gives about -3LUFS. In all cases "LU" is used for loudness relative to some target loudness.

The biggest difference is that A/85 uses a target loudness of -24LUFS whereas R128 uses a target loudness of -23LUFS. You can select either (or another value) using the target level control and the LU values shown will change to correspond.

7.2.4.3 Meters

The three main meters are labelled "M", "S" and "I". These all measure the same underlying loudness but window the results in different ways:

- Momentary Loudness ("M") averages results over a 400ms window. This meter is the quickest to respond to changes in loudness.
- Short Term Loudness ("S") averages results over a 3sec window.
- Integrated Loudness ("I") averages momentary results since the stream started or was last reset. Multistage gating (-70LUFS absolute and -10dB relative) is used to exclude quiet periods from the results, as described in BS.1770.

The history graph shows the last five minutes of the overall level, displaying the range from minimum to maximum for each second.

7.2.5 Controls

7.2.5.1 Control: Scale Range

As recommended in BS.1770 and elsewhere, two scales are provided:

- The "Type 1" range is from -18LU to +9LU.
- The "Type 2" range is from -36LU to +18LU.

Note that this range is independent of the scale units provided, but that the definition of "LU" depends on the target level.

7.2.5.2 Control: Scale Units

The scale values (shown between the history and meters) can be shown in LUFS units (absolute loudness) or LU units (relative to the target level). This does not change the way the plots themselves are shown.

7.2.5.3 Control: Target Level

The meter assumes that you are targeting a particular loudness level, specified in LUFS. EBU R128 recommends -23LUFS and ATSC A/85 recommends -24LUFS. Other values are sometimes needed and values between -36 and +6 can be entered into the text box provided.

Actual graphing is performed using LU units. LU values are simply LUFS values minus the target level. As graphing is always performed in LU (regardless of scale units), changing the target level effectively moves the plots up or down.

7.2.5.4 Control: Active

When the "active" button is on (lit), measurement will occur. When it is off, previous results are kept but new results are not collected.

This means that this button can be used to start measurement, pause and continue.

7.2.5.5 Control: Reset

This button completely resets all meters and history. It can be used when measurement is active or not.
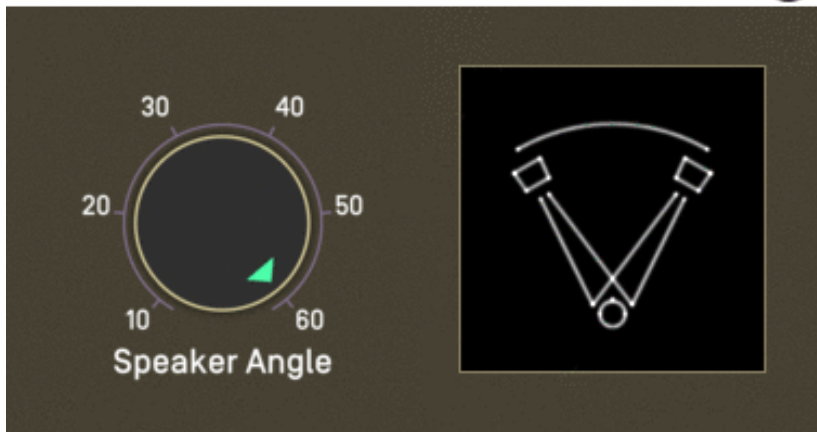
7.2.6 Presets

Presets are available for this plugin:

- EBU Type 1
- EBU Type 2
- ATSC Type 1
- ATSC Type 2

7.3 Stereo Crosstalk Cancellation

7.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

7.3.2 Audio

	Channels	Content
Input	2	Stereo
Output	2	Stereo

7.3.3 Controls

- Angle

7.3.4 Description

When you listen to stereo using normal stereo speakers (rather than headphones), the sound coming from the left speaker reaches both of your ears, not just the left one. Similarly for the right speaker. The sound that travels from the left speaker to the right ear and the sound that travels from the right speaker to the left ear is known as "crosstalk".

This plugin modifies a stereo signal to cancel out some of this effect, although this only works well under highly controlled conditions. This is known as "crosstalk cancellation" (CTC).

7.3.4.1 Why would I want to do that?

Mostly, you would not. Usually, stereo material is intended to be played on stereo speakers and this crosstalk is normal.

However, it can be useful when you have binaural 3D stereo material that is intended to be listened to on headphones (where crosstalk does not occur naturally) and you want to play it on stereo speakers. With this processing it is possible to produce 3D sound using just two speakers.

Binaural 3D stereo recordings can be made with microphones mounted in real or modelled heads, or by software such as Blue Ripple Sound's Binaural Surround processors or O3A Decoder for

Headphones.

7.3.5 Severe Warnings!

There are things about this plugin that you should be aware of:

- Unlike ambisonic decoding methods, this sort of crosstalk cancellation processing is *extremely* sensitive to listener location. It has a *tiny* sweet spot and a head rotation or a head movement of a few centimetres can be enough to ruin the image. The tiny sweet spot also means that this technique normally will not work for more than one listener at once.
- Errors in speaker angle can ruin the image.
- Acoustically reflective surfaces can ruin the image.
- When the image does not work for any reason, the processing used here can result in strong sound colouration.

So, if you are considering using this effect on material that may be hear somewhere that is not set up correctly, be really careful to listen to the results from different angles and ideally on different speakers and on headphones. It's a great effect when it works, but setting everything up correctly can be difficult and a lot can go wrong.
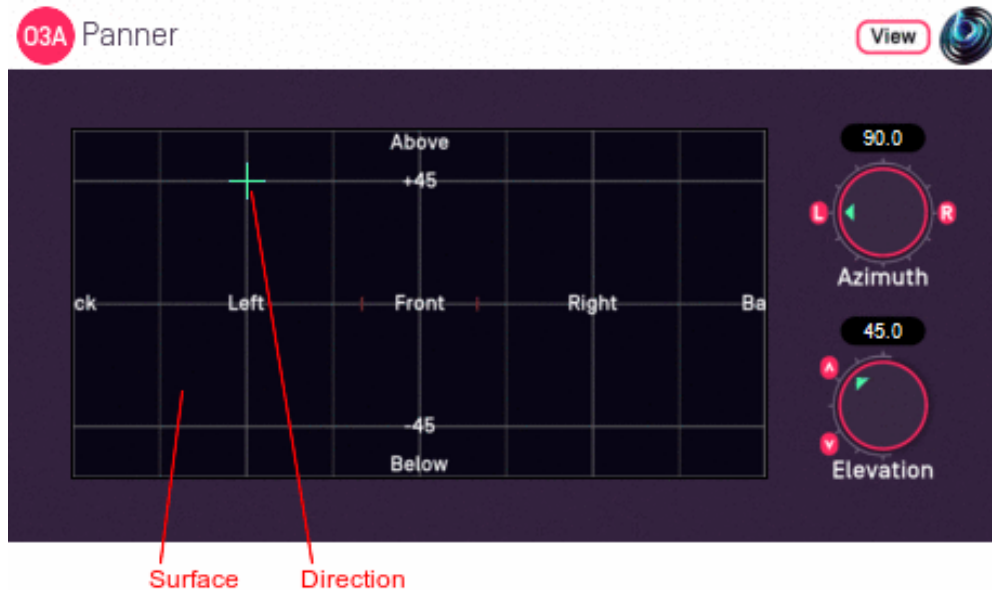
7.3.6 Controls

7.3.6.1 Control: Angle

This plugin has just one control, the angle between the two stereo speakers. For a normal stereo set-up this should be 60 degrees. Desktop and laptop computer speakers tend to be narrower than this, typically 20 or 30 degrees.

Values from 10 to 60 degrees can be selected, in 5 degree steps.

8 O3A Core - Panners

8.1 O3A Panner

8.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.1.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

8.1.3 Controls

- Azimuth
- Elevation

8.1.4 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.

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.
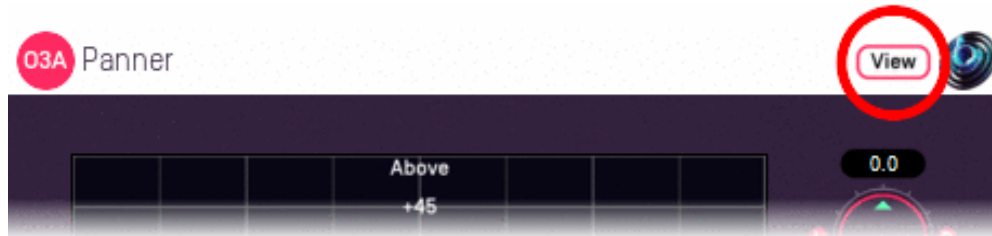
When this is pressed, the plugin is connected to the View or ViewVR application and is displayed there.

View Support	ViewVR Support
The current direction is shown using a cross.	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.

8.1.6 Controls
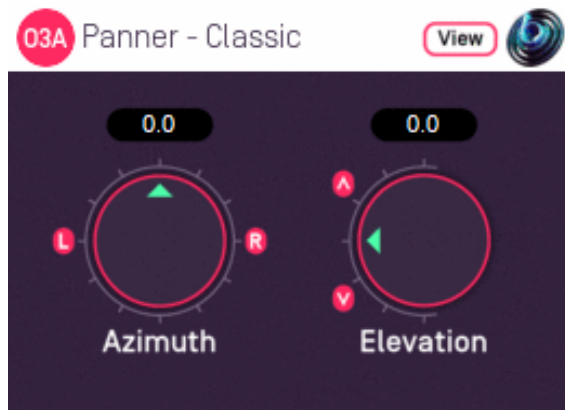
8.1.6.1 Controls: Azimuth and Elevation

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.

8.2 O3A Panner - Classic

8.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.2.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

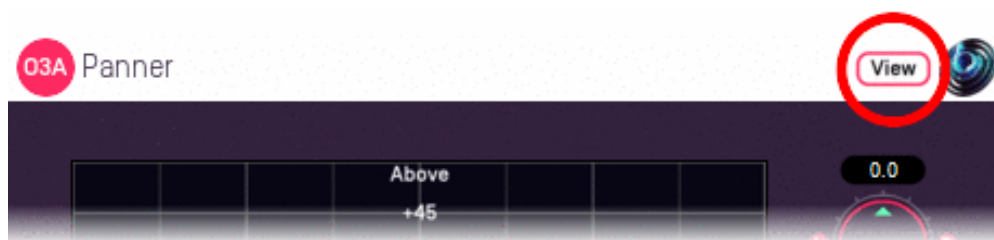### 8.2.3 Controls

- Azimuth
- Elevation

8.2.4 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.

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
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.

8.2.6 Controls

8.2.6.1 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.

8.3 O3A Panner - Hemisphere
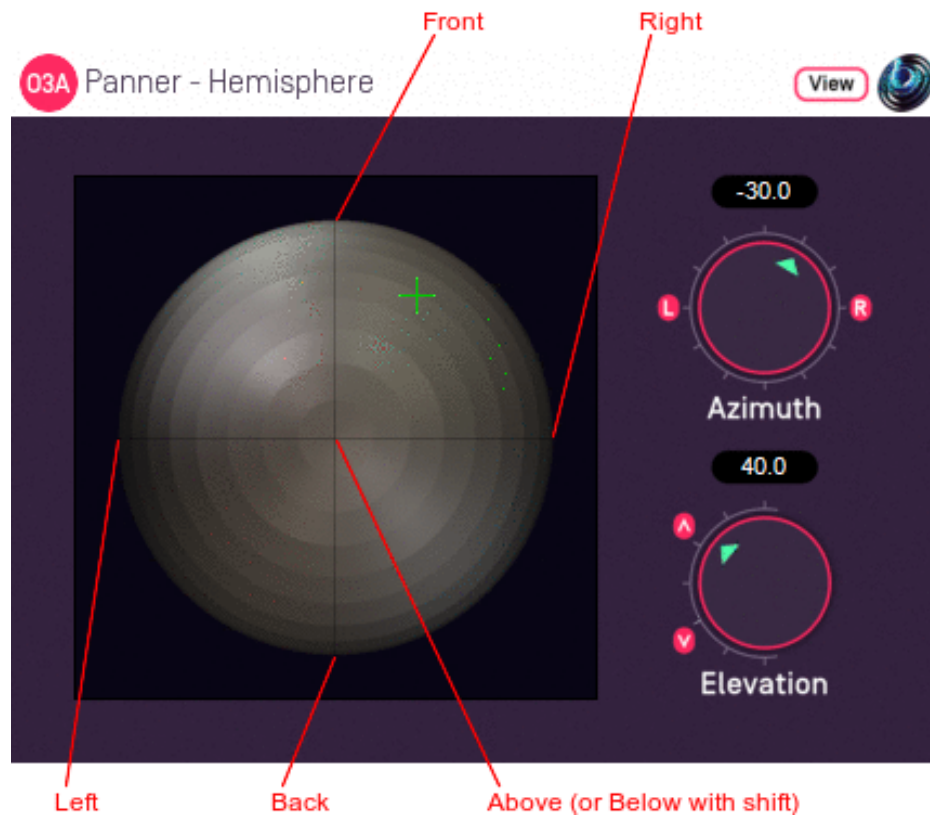
8.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.3.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

8.3.3 Controls

- Azimuth
- Elevation

8.3.4 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:
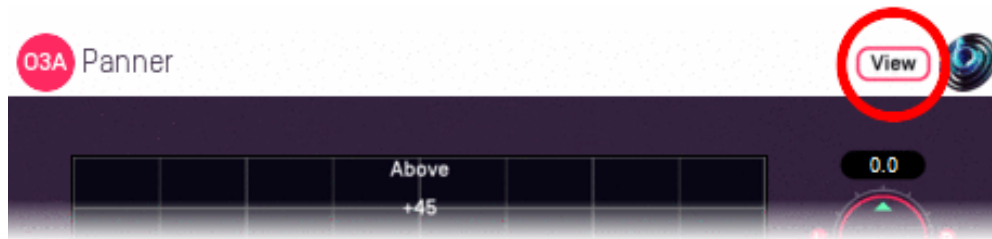
- Use the "shift" key to select the lower hemisphere (so sounds are below rather than above).
- Use the "alt" key to lock the elevation (so sounds move in a horizontal circle).

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

8.3.5 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.	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.

8.3.6 Controls

8.3.6.1 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.

8.4 O3A Panner - Large
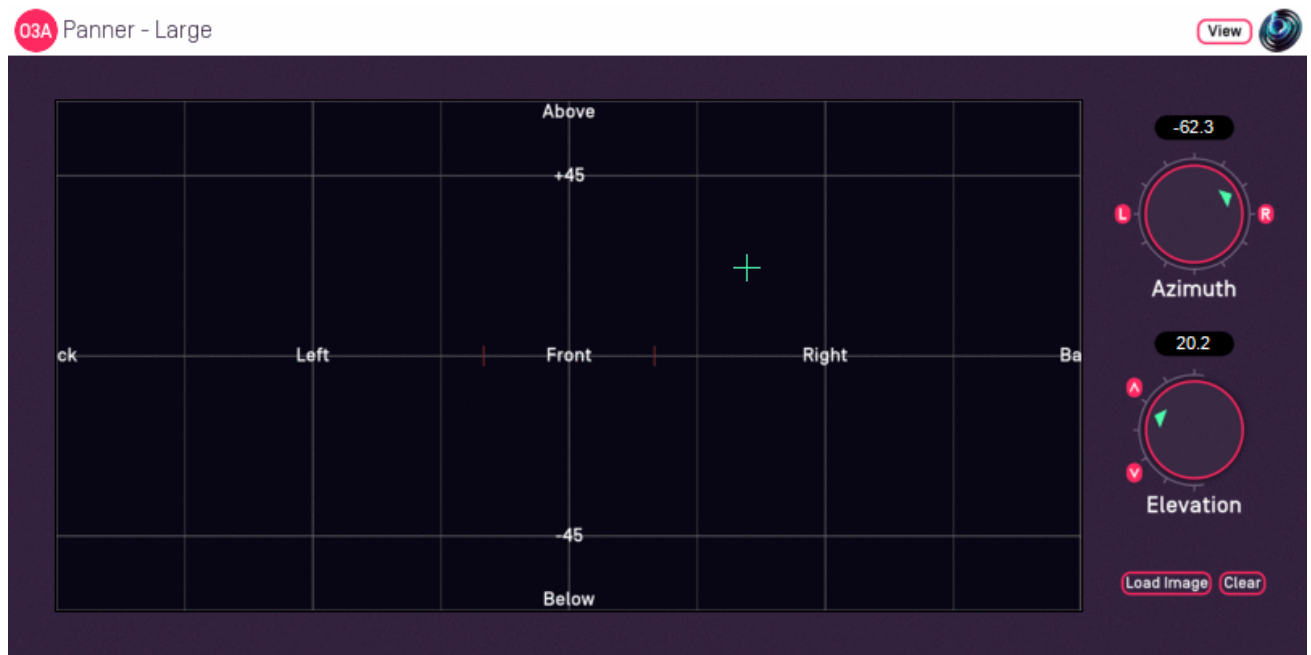
8.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.4.2 Audio

	Channels	Content
Input	1	Mono
Output	16	O3A

8.4.3 Controls

- Azimuth
- Elevation

8.4.4 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.

8.4.4.1 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.
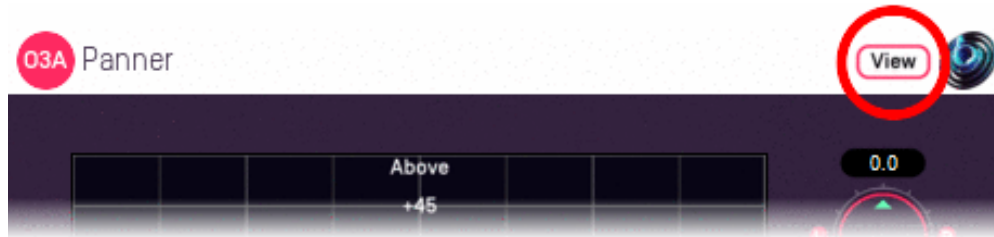
The target image size is 636x316. Other sizes are resized to fit.

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

8.4.5 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.	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.

8.4.6 Controls

8.4.6.1 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.

8.5 O3A Panner - Two Channel
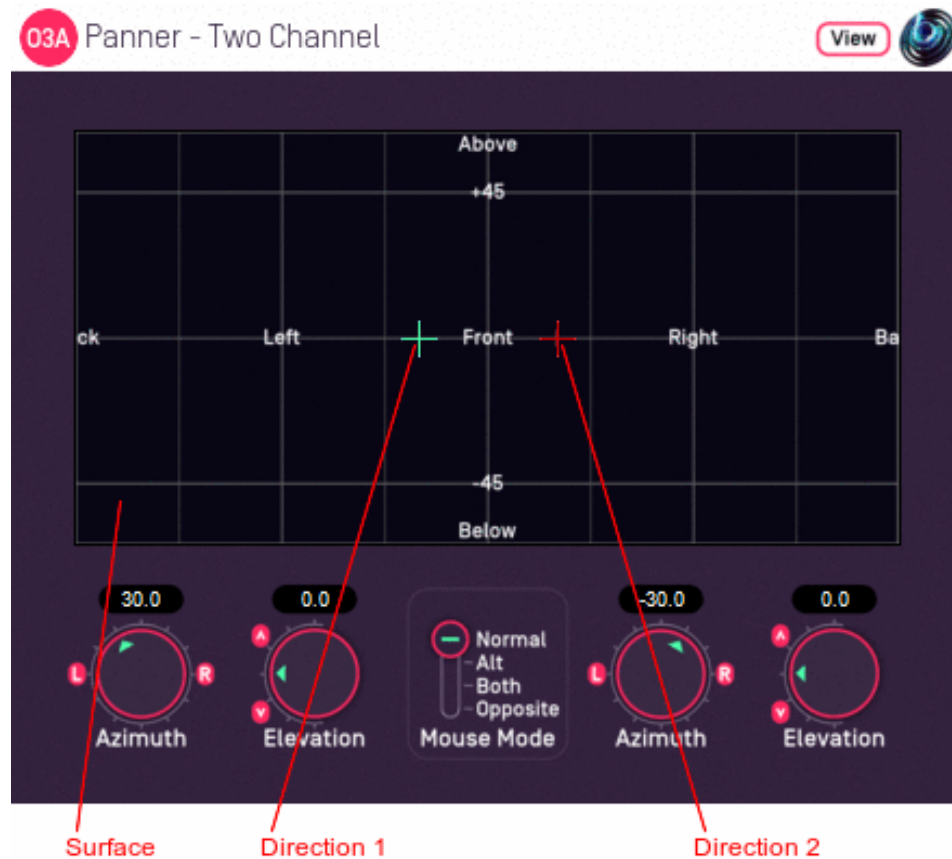
8.5.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

8.5.2 Audio

	Channels	Content
Input	2	Stereo
Output	16	O3A

8.5.3 Controls

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

8.5.4 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.
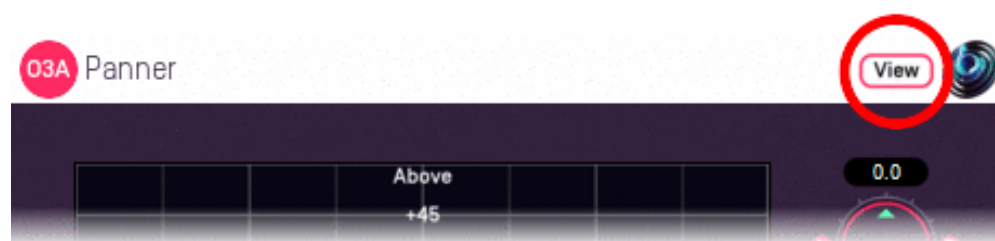
Mouse Mode	Key(s) Held	Function
Norm		Normal operation. Left click to move the first source, right click to move the second source.
Alt	Alt	Flips the source moved when the control surface is clicked, so a left click moves the second source rather than the first (and vice versa). Typically used to select 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.

8.5.5 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
The current panner directions are shown using green/blue and red crosses, similarly to the plugin display. Point at one and hold the left mouse button to select it. Then, drag to set its azimuth and elevation.	The current panner directions are shown using green/blue and red beams and crosses, similarly to the plugin display. 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.

8.5.6 Controls

8.5.6.1 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).

8.5.6.2 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.

8.6 O3A Panner - Eight Channel
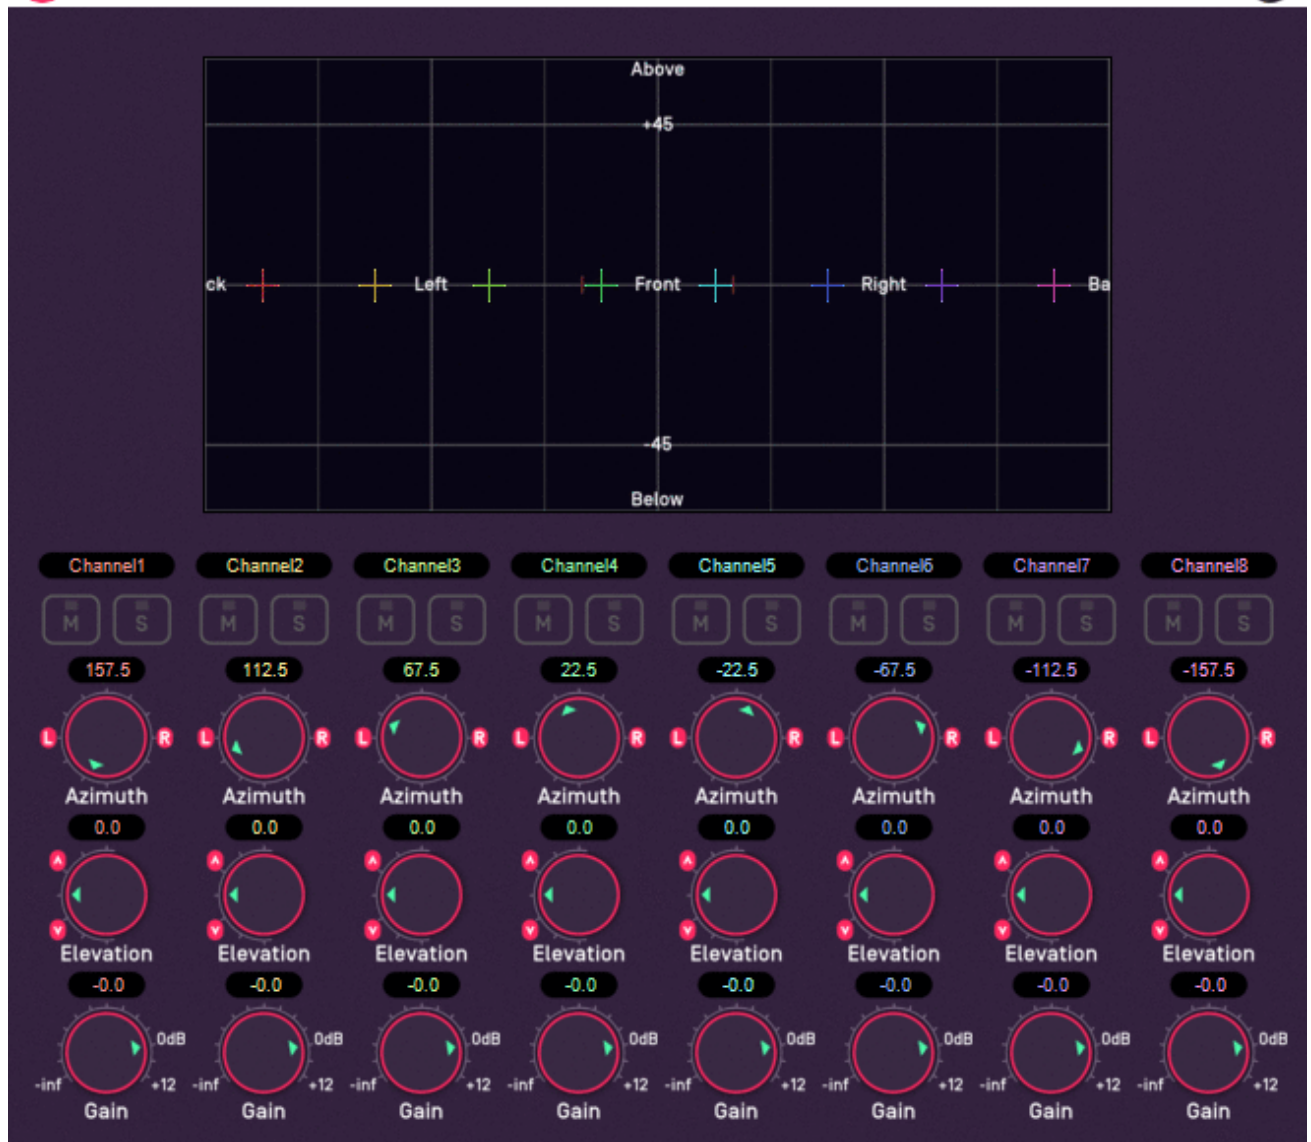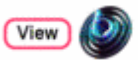
8.6.1 Host Support

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

8.6.2 Audio

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

8.6.3 Controls

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

- Mute 1 - 8
- Solo 1 - 8

8.6.4 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.

8.6.4.1 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.

8.6.4.2 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.

8.6.5 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
The current panner directions are shown using colour-coded crosses and labels, similarly to the plugin display. Point at one and hold the left mouse button to select it. Then, drag to set its azimuth and elevation.	The current panner directions are shown using colour-coded beams, crosses and labels, similarly to the plugin display. 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.

8.6.6 Controls

8.6.6.1 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.

8.6.6.2 Controls: Gain 1 - 8

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

8.6.6.3 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.
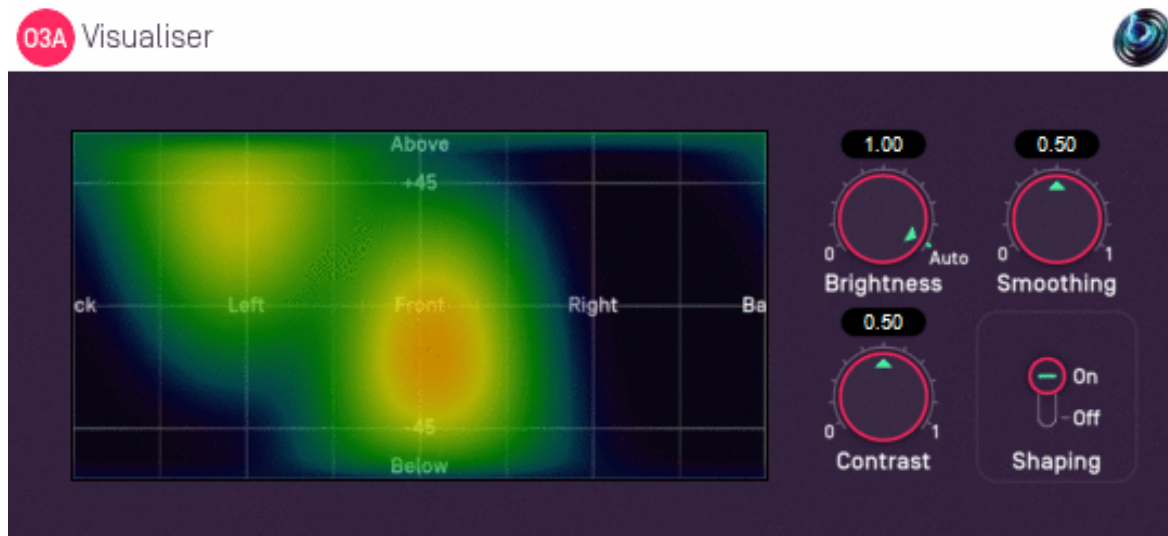
8.6.7 Presets

Presets are available for this plugin:

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

9 O3A Core - Visualisers

9.1 O3A Visualiser

9.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

9.1.3 Controls

- Brightness
- Contrast
- Smoothing
- Shaping

9.1.4 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.

9.1.5 Controls

9.1.5.1 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.

9.1.5.2 Control: Contrast

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

9.1.5.3 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.

9.1.5.4 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.

9.2 O3A Visualiser - Hemisphere
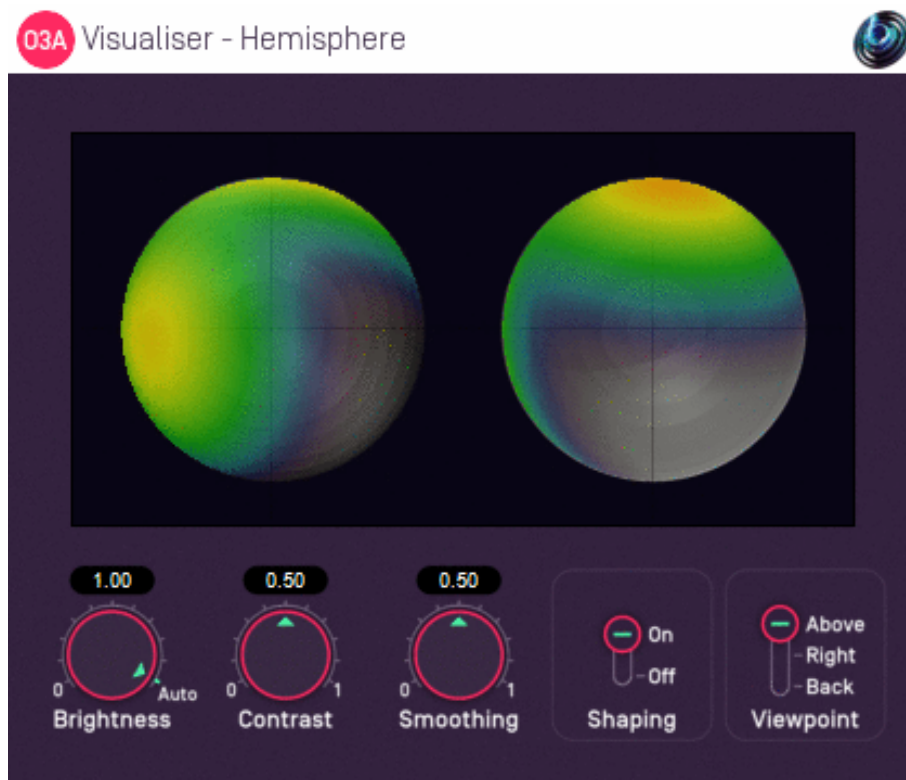
9.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

9.2.3 Controls

- Brightness
- Contrast
- Smoothing
- Shaping
- Viewpoint

9.2.4 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.

9.2.5 Controls

9.2.5.1 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.

9.2.5.2 Control: Contrast

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

9.2.5.3 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.

9.2.5.4 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.

9.2.5.5 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.

9.3 O3A Colourizer
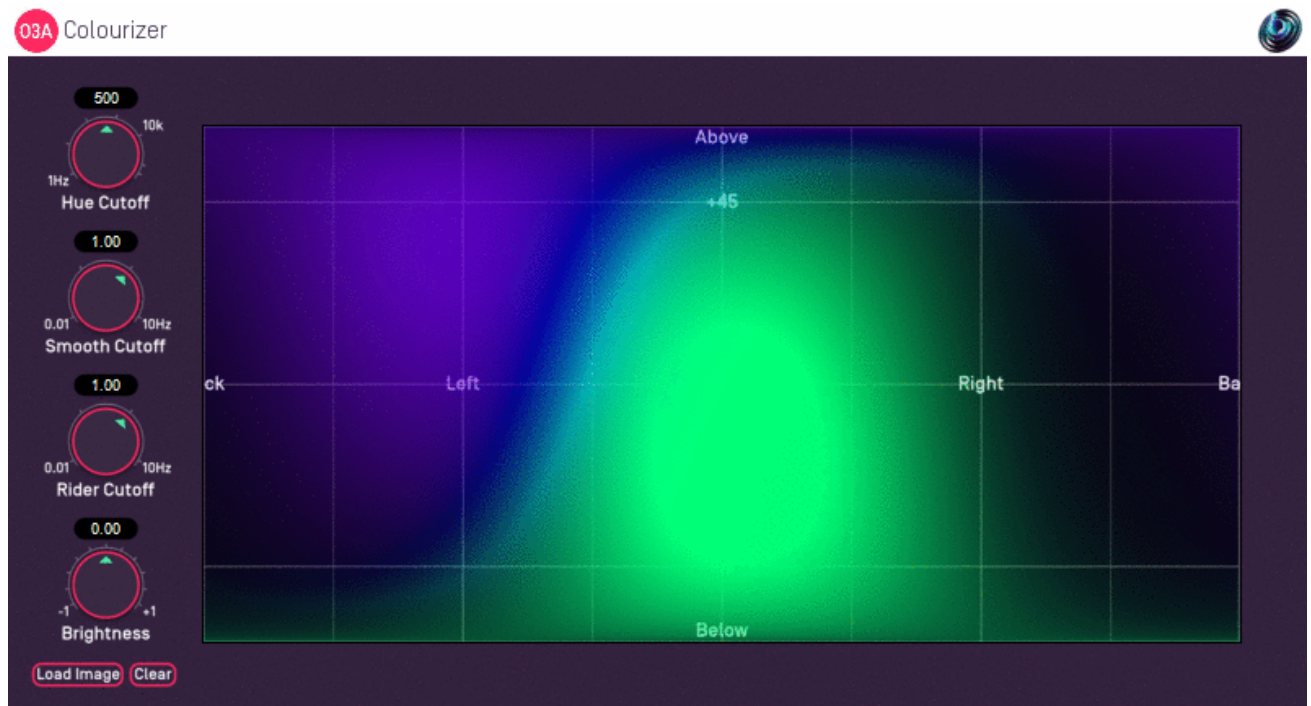
9.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.3.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

9.3.3 Controls

- Hue Cutoff
- Smooth Cutoff
- Rider Cutoff
- Brightness

9.3.4 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.

9.3.4.1 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.

9.3.5 Controls

9.3.5.1 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.

9.3.5.2 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.
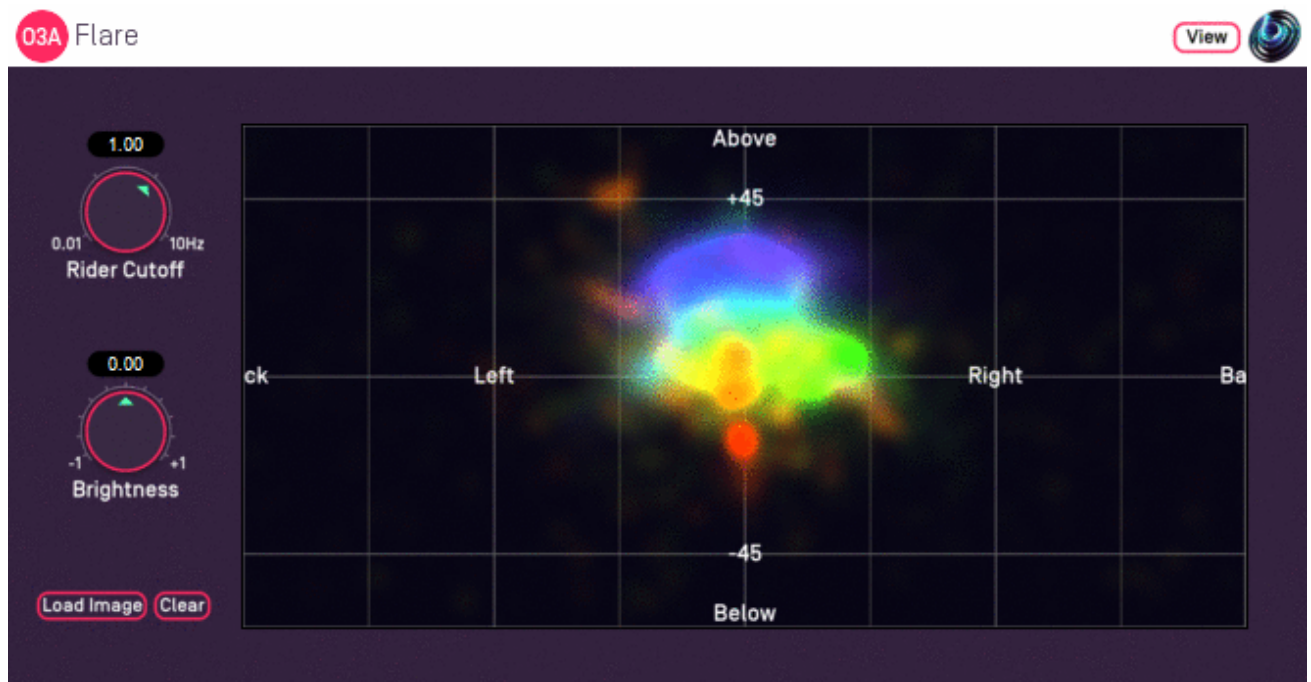
9.3.5.3 Control: Rider Cutoff

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.

9.3.5.4 Control: Brightness

This modifies the overall brightness of the image.

9.4 O3A Flare

9.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

9.4.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

9.4.3 Controls

- Rider Cutoff
- Brightness

9.4.4 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 plugins 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.

9.4.4.1 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.

9.4.5 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.
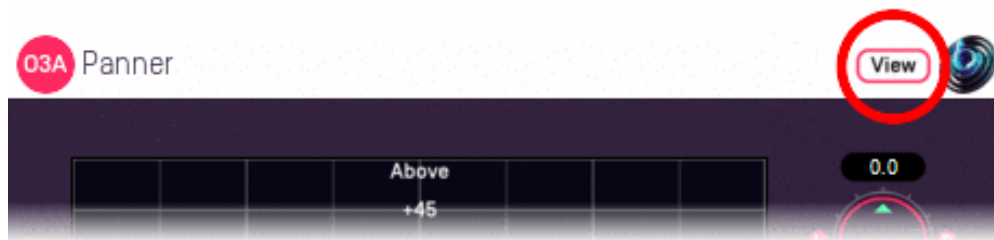
When this is pressed, the plugin is connected to the View or ViewVR application and is displayed there.

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.

9.4.6 Controls

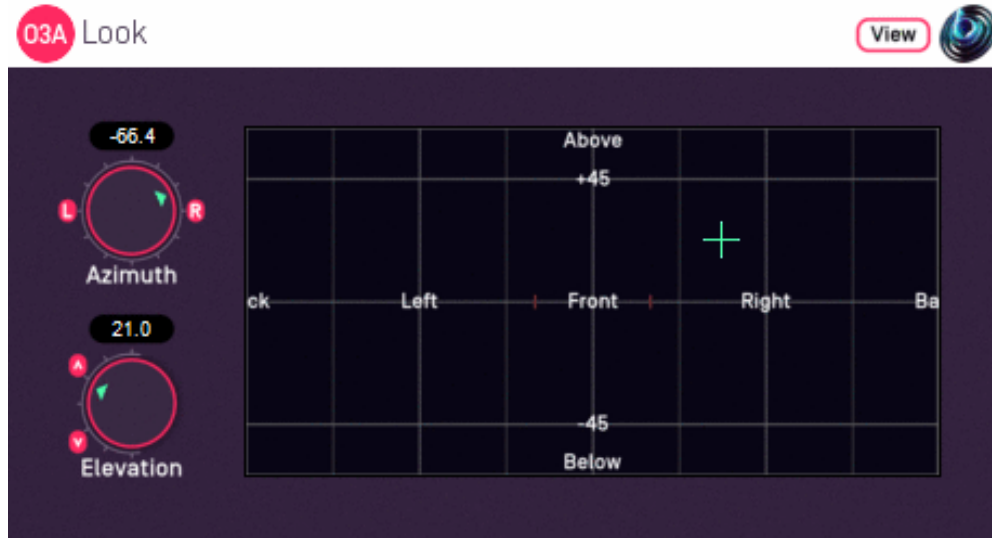**9.4.6.1 Control: Rider Cutoff**

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.

9.4.6.2 Control: Brightness

This modifies the overall brightness of the image.

10 O3A Core - Rotation

10.1 O3A Look

10.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

10.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

10.1.3 Controls

- Azimuth
- Elevation

10.1.4 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).

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. 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.

10.1.6 Controls

10.1.6.1 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.

10.2 O3A Rotation

10.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

10.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

10.2.3 Controls

- Yaw
- Pitch
- Roll
- Ordering

10.2.4 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.

10.2.5 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.

10.2.6 Controls

10.2.6.1 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.

10.2.6.2 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).
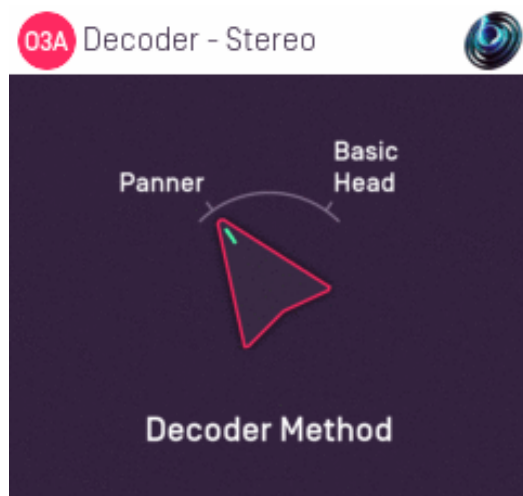
When simulating changes in the listener's orientation, we generally recommend the Y-P-R ordering. Then, the three dials can be thought of as turning the listener's head horizontally, then pitching it

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

11 O3A Core - Decoders and Converters

11.1 O3A Decoder - Stereo

11.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

11.1.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Stereo

11.1.3 Controls

- Method

11.1.4 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.

11.1.5 Controls

11.1.5.1 Control: Method

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

Method	Description
Panner	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. 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.
Basic Head	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. **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. 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. 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.

11.1.5.2 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.

11.2 O3A Decoder - Mono

11.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

11.2.2 Audio

	Channels	Content
Input	16	O3A
Output	1	Mono

11.2.3 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.

11.3 O3A Decoder - 5.1 Basic

11.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

11.3.2 Audio

	Channels	Content
Input	16	O3A
Output	6	5.1 Surround

11.3.3 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.

11.3.3.1 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.

11.4 O3A Decoder - FuMa
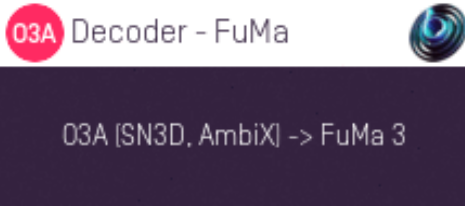
11.4.1 Host Support

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

11.4.2 Audio

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

11.4.3 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.

11.4.3.1 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.

11.5 O3A Injector - FuMa
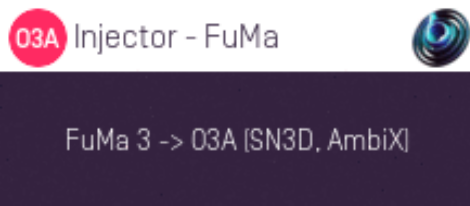
11.5.1 Host Support

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

11.5.2 Audio

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

11.5.3 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.

11.5.3.1 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.

11.6 O3A Decoder - O1A and O2A

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

11.6.1 Host Support

11.6.1.1 O3A Decoder - O1A

Host Type	Support
AAX	Yes
VST2	Yes

11.6.1.2 O3A Decoder - O2A

Host Type	Support
AAX	Yes
VST2	Yes

11.6.2 Audio

11.6.2.1 O3A Decoder - O1A

	Channels	Content
Input	16	O3A
Output	4	O1A

11.6.2.2 O3A Decoder - O2A

	Channels	Content
Input	16	O3A
Output	9	O2A

11.6.3 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.

11.7 O3A Injector - O1A and O2A

11.7.1 Host Support

11.7.1.1 O3A Injector - O1A

Host Type	Support
AAX	Yes
VST2	Yes

11.7.1.2 O3A Injector - O2A

Host Type	Support
AAX	Yes
VST2	Yes

11.7.2 Audio

11.7.2.1 O3A Injector - O1A

	Channels	Content
Input	4	O1A
Output	16	O3A

11.7.2.2 O3A Injector - O2A

	Channels	Content
Input	9	O2A
Output	16	O3A

11.7.3 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.

12 O3A Core - Virtual Microphones

12.1 O3A Beamer
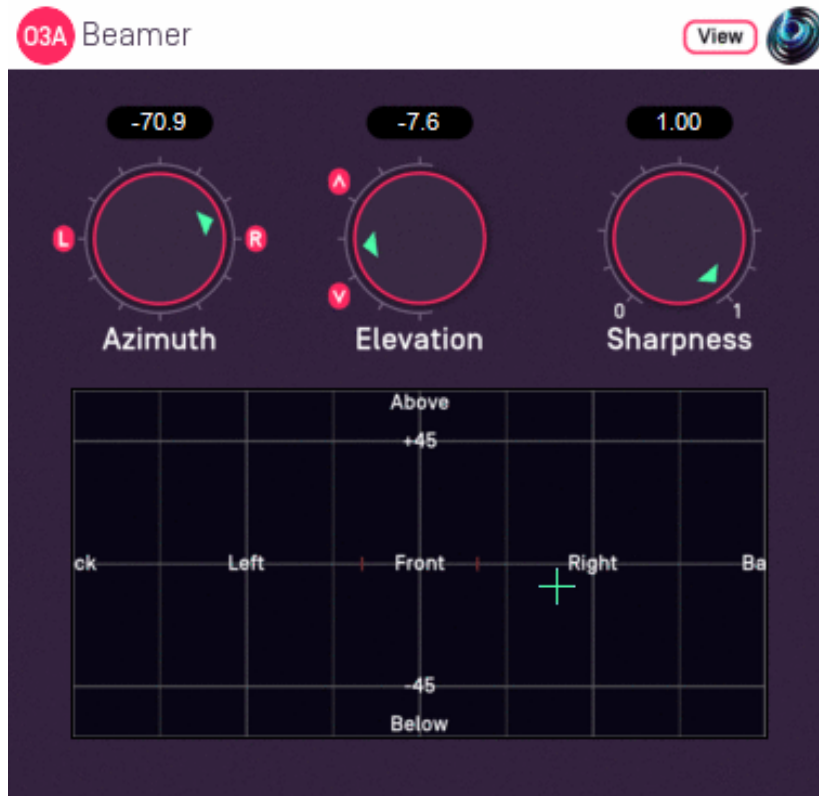
12.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

12.1.2 Audio

	Channels	Content
Input	16	O3A
Output	1	Mono

12.1.3 Controls

- Azimuth
- Elevation
- Sharpness

12.1.4 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.

12.1.5 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.
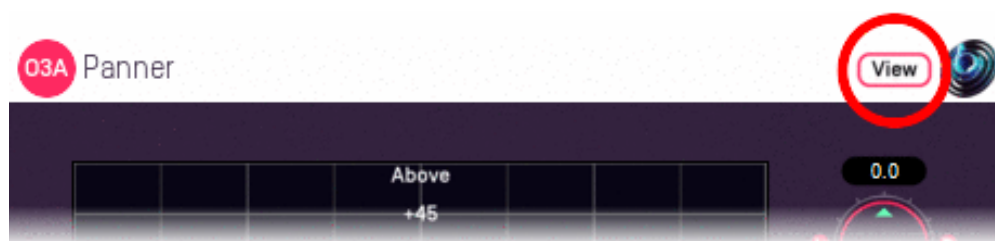
When this is pressed, the plugin is connected to the View or ViewVR application and is displayed there.

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.

12.1.6 Controls

12.1.6.1 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.

12.1.6.2 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.

12.2 O3A Virtual Microphone
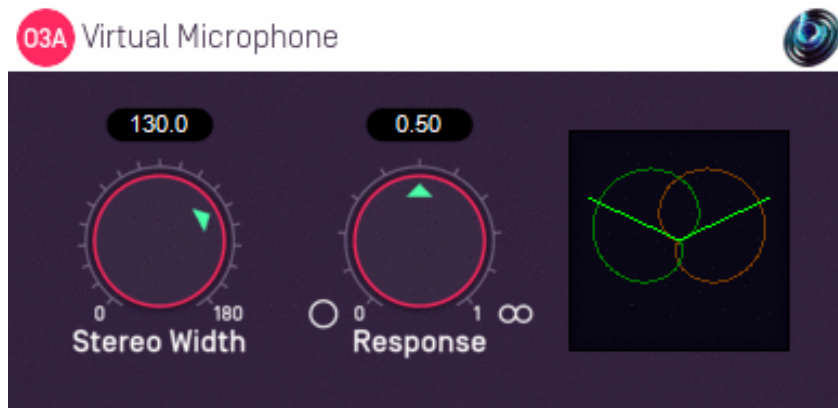
12.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

12.2.2 Audio

	Channels	Content
Input	16	O3A
Output	2	Stereo

12.2.3 Controls

- Stereo Width
- Response

12.2.4 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.

12.2.5 Controls

12.2.5.1 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.

12.2.5.2 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

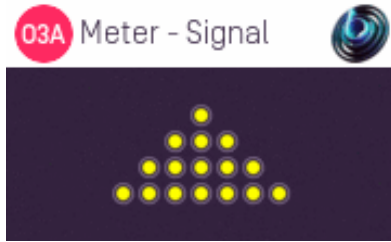For instance, 1.00 (figure-of-eight) should be used for Blumlein X-Y.

12.2.6 Presets

Presets are available for this plugin:

- Cardioids 130
- Cardioids 90
- Blumlein X/Y

13 O3A Core - Metering

13.1 O3A Meter - Signal

13.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

13.1.2 Audio

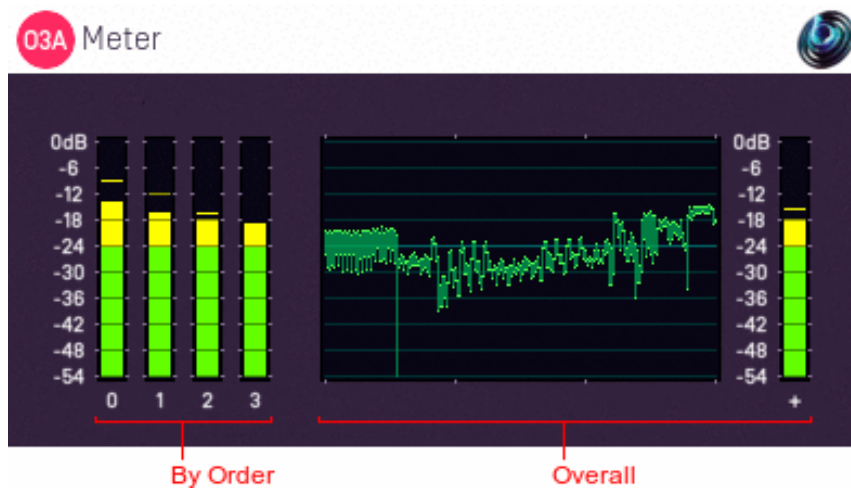|	Channels	Content
Input	16	O3A
Output	16	O3A

13.1.3 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.

13.1.3.1 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.

13.2 O3A Meter

13.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

13.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

13.2.3 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.

13.2.3.1 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
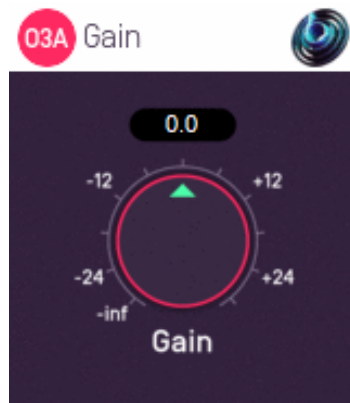
Decoding library, which will produce comparable numbers when set to use LUFS units, subject to directional and loudness weighting.

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

14 O3A Core - Soundfield Modification

14.1 O3A Gain

14.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

14.1.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

14.1.3 Controls

- Gain

14.1.4 Description

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

14.1.5 Controls

14.1.5.1 Control: Gain

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

14.2 O3A Spatial Delay

14.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

14.2.2 Audio

	Channels	Content
Input	16	O3A
Output	16	O3A

14.2.3 Controls

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

14.2.4 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*
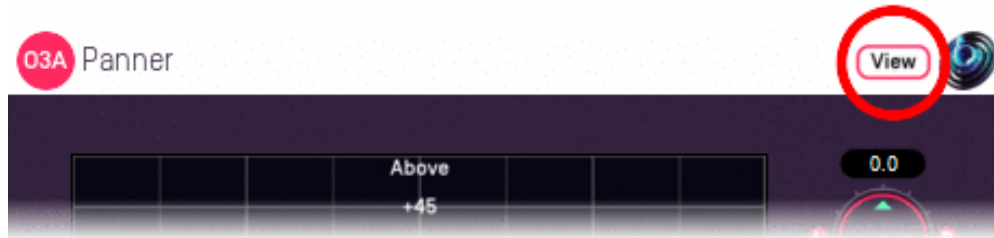
the feedback loop. The wet gain does not, applying only to the final output.

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

14.2.5 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.

14.2.6 Controls

14.2.6.1 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.

14.2.6.2 Controls: Dry Gain and Wet Gain

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

14.2.6.3 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.

14.2.6.4 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.

14.2.6.5 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.

14.2.6.6 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.

14.2.6.7 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).

14.2.6.8 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.

15 Appendix: O3A Streams

15.1 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.

15.2 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.

15.3 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.

15.4 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.

15.5 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.