O7A Decoding Plugins

v2.4.0

Table of Contents

1 Introduction

The O7A Decoding plugins decode (or "render") O7A to various other formats including 5.1, 7.1.2 and binaural stereo.

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

1.1 Compatibility

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

1.1.1 AAX

Most of the O7A plugins exist as AAX versions which are compatible with Pro Tools Ultimate version 2023.6 or later 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 O7A plugins use internal buffering with a length of 128 samples. For smooth CPU load, you may wish to ensure your DAW buffer size is a multiple of this.

1.2 Getting Started

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

1.3 The License Manager

The Blue Ripple Sound License Manager application can be used to move license keys around between computers.

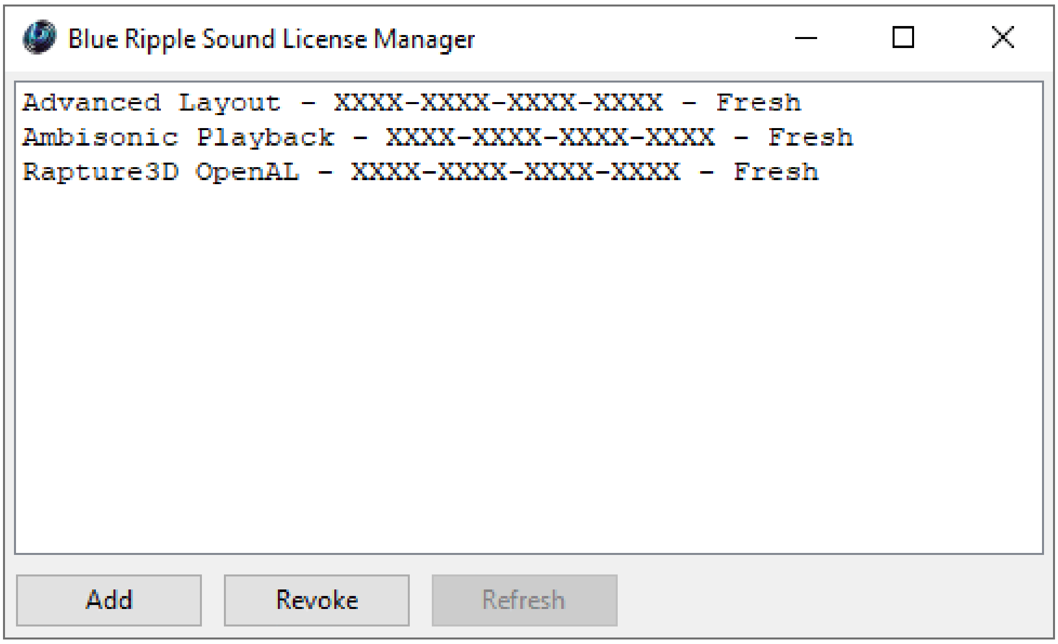

The application is generally available in the Start Menu on Windows, and in your installation directory, which is typically:

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

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

1.4 Technical Requirements

1.4.1 Operating System

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

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

1.4.2 VST Host

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

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

1.4.3 AAX Host

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

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

1.4.4 PC Hardware

Please check your PC meets the following requirements:

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

1.4.5 Permissions

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

1.4.6 Internet Connection Required

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

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

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

2 Decoding

2.1 When should you decode?

Ideally, decoding should only ever happen in the final listening space. If the O7A stream (or a lower order version) is sent there it can be decoded with the best algorithm for that particular space and speaker layout. Or, for Virtual Reality, the O7A 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 O7A 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 O7A 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 O7A streams.

2.2 Choosing a Decoder

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

2.2.1 Headphones

If you are listening on headphones, the binaural decoder O7A Decoder - Headphones will almost certainly work best, although you could also try the basic Stereo decoder (from the O7A Core) for a simpler mix which may play more naturally on ordinary stereo speakers.

2.2.2 Stereo Speakers

The simplest and most robust option here is the basic Stereo decoder from the O7A Core.

An alternative to decoding to stereo in this way is to use a virtual stereo microphone (also from the O7A 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.

2.2.3 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.

2.2.4 Other Ambisonic Formats

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

Support for "FuMa" is available in the O7A Core.

To output in the "N3D" format (using "ACN" channel ordering), convert using the O7A Decoder - N3D plugin.

To reduce the order, simply use the first four channels (first order), nine channels (second order), sixteen channels (third order) and so on, and discard the rest.

2.3 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.

2.4 Decoder Levels and Loudness

All O7A decoders and custom decoders generated using Blue Ripple Sound's Rapture3D "Advanced" edition are designed to produce similar overall output levels.

The O7A Meter or O7A Meter - Karma plugins can be used to estimate these levels from the O7A 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 O7A stream to a file, you may wish to use the O7A Brickwall limiter to protect against clipping.

3 Binaural Decoding

3.1 O7A Decoder - Headphones

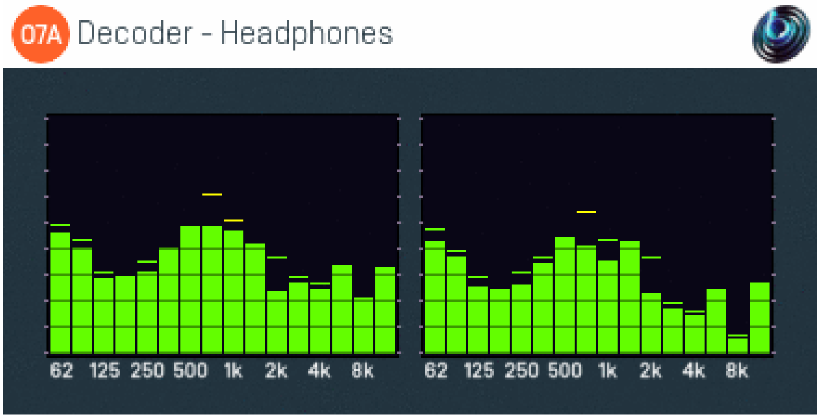

3.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

3.1.2 Audio

	Channels	Content
Input	64	O7A
Output	2	Binaural Stereo

3.1.3 Description

This plugin takes a seventh order ambisonics (O7A) 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 O7A 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.

3.1.3.1 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 seventh order (O7A) version of the decoder.

4 Loudspeaker Decoding

4.1 O7A Decoder - Quad

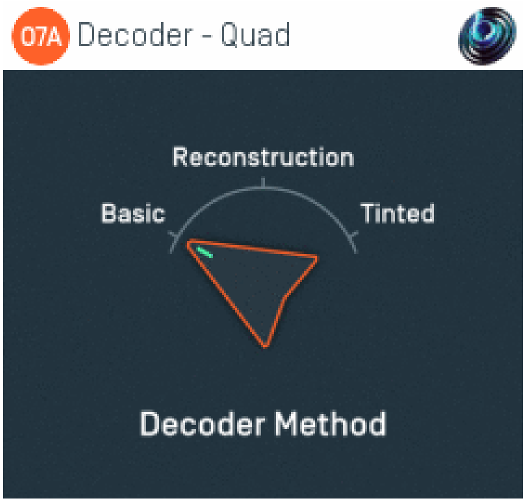

4.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.1.2 Audio

	Channels	Content
Input	64	O7A
Output	4	Quad

4.1.3 Controls

- Method

4.1.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as speaker feeds for four speakers set out at the corners of a square.

4.1.4.1 Channels

The channel ordering used is:

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

4.1.5 Controls

4.1.5.1 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.

4.2 O7A Decoder - 5.1

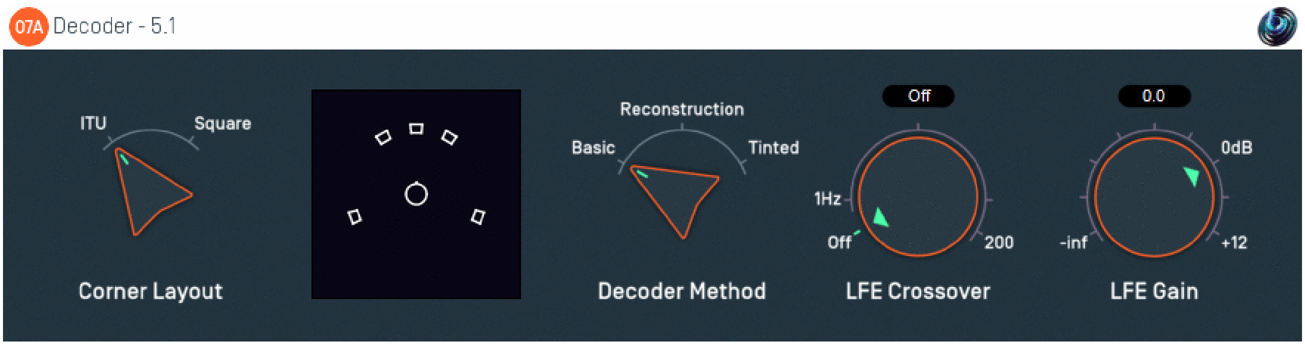

4.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.2.2 Audio

	Channels	Content
Input	64	O7A
Output	6	5.1 Surround

4.2.3 Controls

- Corners
- Method
- LFE Crossover
- LFE Gain

4.2.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as a 5.1 surround mix.

4.2.4.1 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

4.2.5 Controls

4.2.5.1 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.

4.2.5.2 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.

4.2.5.3 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.

4.2.5.4 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.

4.3 O7A Decoder - 7.1

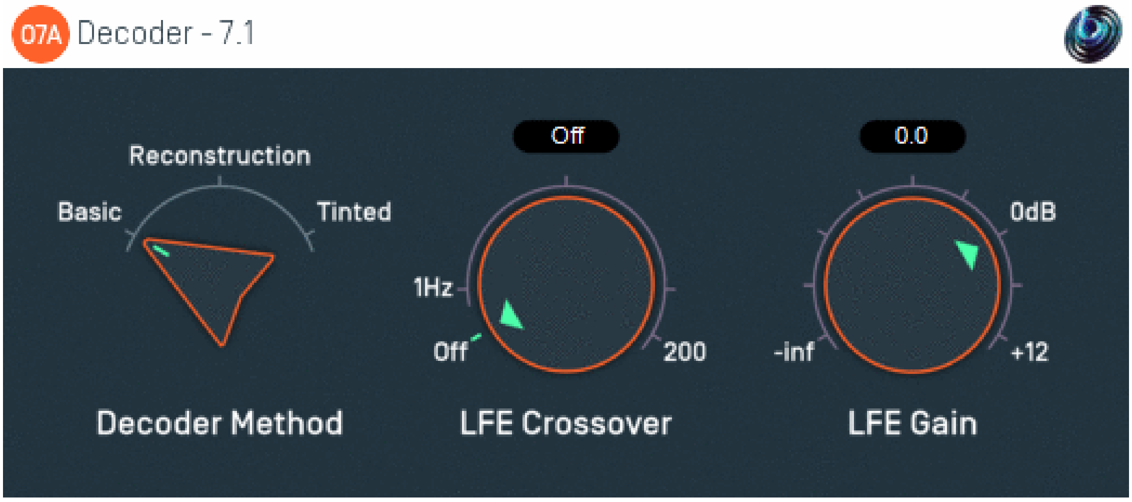

4.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.3.2 Audio

	Channels	Content
Input	64	O7A
Output	8	7.1 Surround

4.3.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.3.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as a 7.1 surround mix.

4.3.4.1 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

4.3.5 Controls

4.3.5.1 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.

4.3.5.2 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.

4.3.5.3 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.

4.4 O7A Decoder - 7.1.2 (Dolby Atmos)

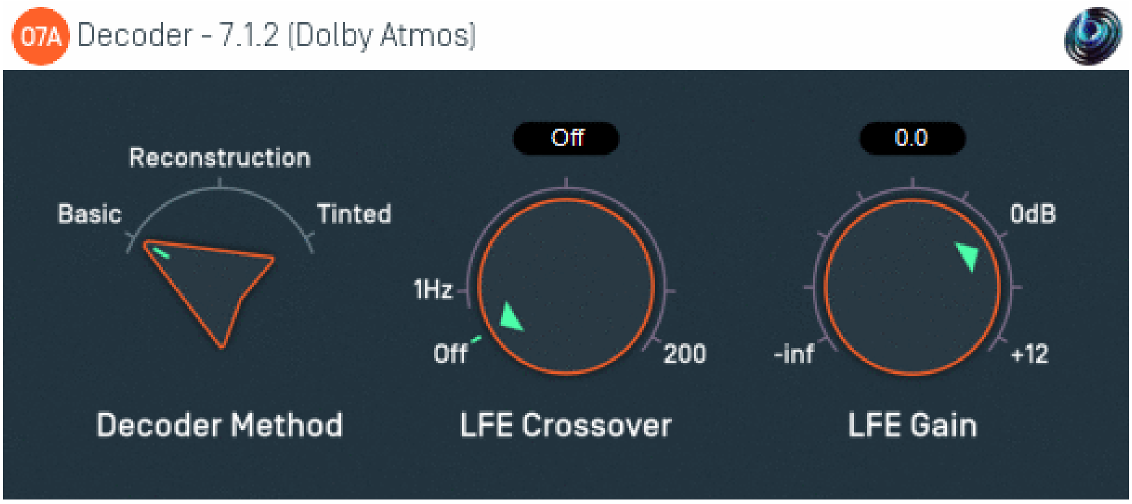

4.4.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.4.2 Audio

	Channels	Content
Input	64	O7A
Output	10	7.1.2 (Dolby Atmos)

4.4.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.4.4 Description

This plugin takes a seventh order ambisonics (O7A) 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.

4.4.4.1 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.

4.4.5 Controls

4.4.5.1 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.

4.4.5.2 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.

4.4.5.3 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.

4.5 O7A Decoder - 7.1.4

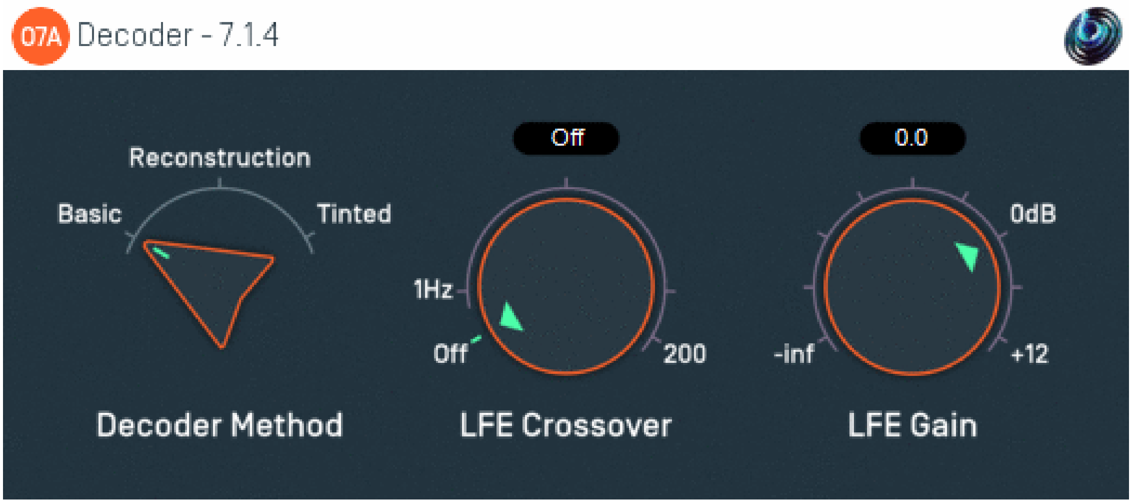

4.5.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

4.5.2 Audio

	Channels	Content
Input	64	O7A
Output	12	7.1.4 Surround

4.5.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.5.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it to 7.1.4.

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.

4.5.4.1 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

4.5.5 Controls

4.5.5.1 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.

4.5.5.2 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.

4.5.5.3 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.

4.6 O7A Decoder - Auro-3D 9.1, 10.1, 11.1 and 13.1

4.6.1 Host Support

4.6.1.1 O7A Decoder - Auro-3D 9.1

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

4.6.1.2 O7A Decoder - Auro-3D 10.1

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

4.6.1.3 O7A Decoder - Auro-3D 11.1

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

4.6.1.4 O7A Decoder - Auro-3D 13.1

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

4.6.2 Audio

4.6.2.1 O7A Decoder - Auro-3D 9.1

	Channels	Content
Input	64	O7A
Output	10	Auro-3D 9.1 (mapped to 5.1.4 for AAX)

4.6.2.2 O7A Decoder - Auro-3D 10.1

	Channels	Content
Input	64	O7A
Output	11	Auro-3D 10.1 (mapped to 7.0.4 for AAX)

4.6.2.3 O7A Decoder - Auro-3D 11.1

	Channels	Content
Input	64	O7A
Output	12	Auro-3D 11.1 (mapped to 7.1.4 for AAX)

4.6.2.4 O7A Decoder - Auro-3D 13.1

	Channels	Content
Input	64	O7A
Output	14	Auro-3D 13.1 (mapped to 7.1.6)

4.6.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.6.4 Description

Four decoder plugins are provided, which produce Auro-3D surround mixes from O7A 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.

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.

4.6.4.1 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.

4.6.5 Controls

4.6.5.1 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.

4.6.5.2 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.

4.6.5.3 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.

4.7 O7A Decoder - IMAX 5.0, 6.0 and 12.0

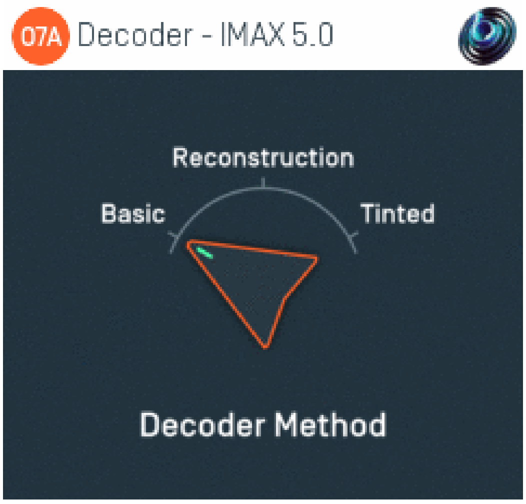

4.7.1 Host Support

4.7.1.1 O7A Decoder - IMAX 5.0

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

4.7.1.2 O7A Decoder - IMAX 6.0

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

4.7.1.3 O7A Decoder - IMAX 12.0

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

4.7.2 Audio

4.7.2.1 O7A Decoder - IMAX 5.0

	Channels	Content
Input	64	O7A
Output	5	IMAX 5.0 (mapped to standard 5.0 for AAX)

4.7.2.2 O7A Decoder - IMAX 6.0

	Channels	Content
Input	64	O7A
Output	6	IMAX 6.0 (mapped to standard 6.0 for AAX)

4.7.2.3 O7A Decoder - IMAX 12.0

	Channels	Content
Input	64	O7A
Output	12	IMAX 12.0 (mapped to 7.1.4 for AAX)

4.7.3 Controls

- Method

4.7.4 Description

These plugins take a seventh order ambisonics (O7A) mix and decode it as IMAX surround mixes.

4.7.4.1 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.

4.7.4.2 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).

4.7.4.3 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.

4.7.5 Controls

4.7.5.1 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.

4.8 O7A Decoder - 22.2

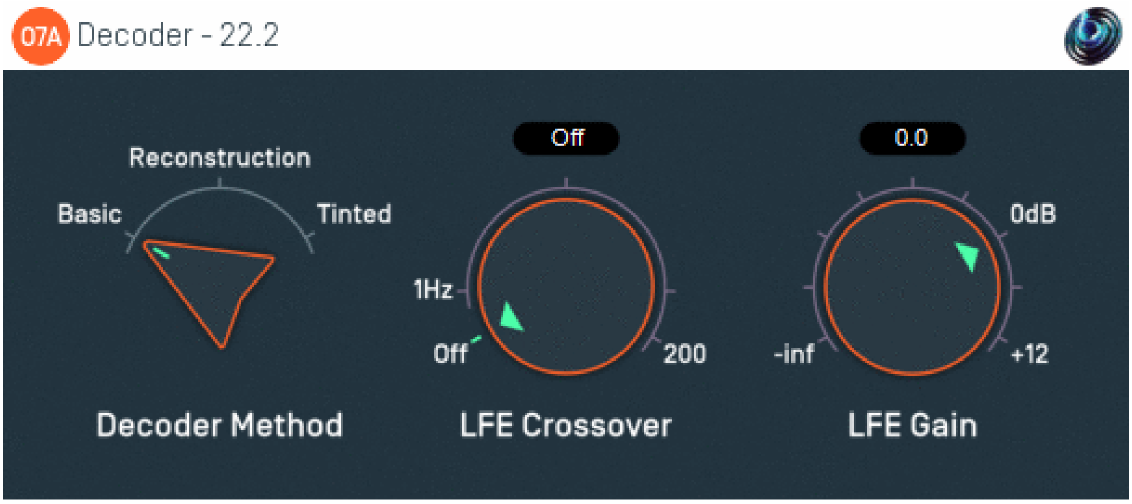

4.8.1 Host Support

Host Type	Support
AAX	No
VST2	Yes

4.8.2 Audio

	Channels	Content
Input	64	O7A
Output	24	Hamasaki 22.2

4.8.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.8.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as a 24 channel Hamasaki 22.2 surround mix.

4.8.4.1 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

4.8.5 Controls

4.8.5.1 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.

4.8.5.2 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.

4.8.5.3 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.

4.9 O7A Decoder - 3D7.1

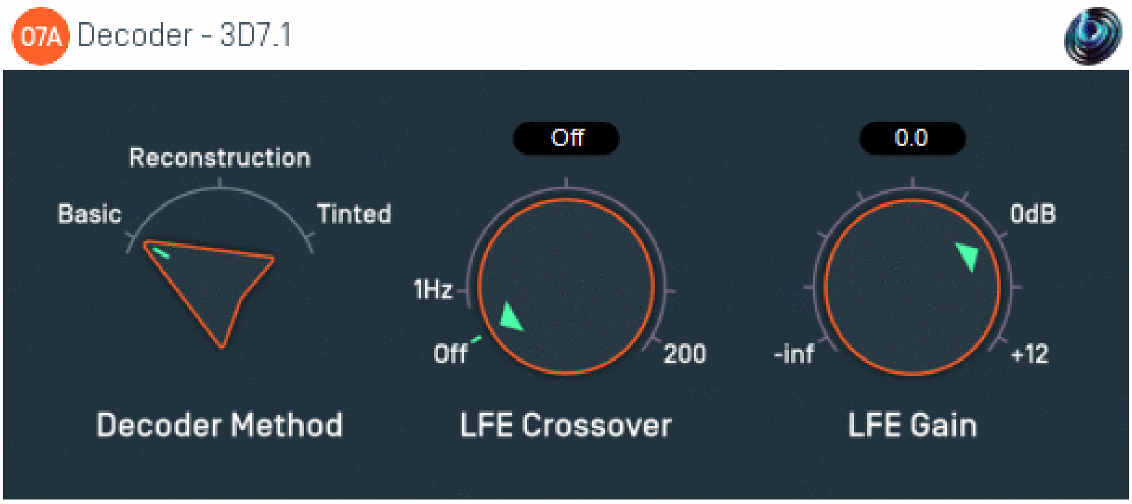

4.9.1 Host Support

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

4.9.2 Audio

	Channels	Content
Input	64	O7A
Output	8	3D7.1 Surround (mapped to 7.1 for AAX)

4.9.3 Controls

- Method
- LFE Crossover
- LFE Gain

4.9.4 Description

This plugin takes a seventh order ambisonics (O7A) 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 O7A decoder for 7.1.

4.9.4.1 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

4.9.4.2 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.

4.9.5 Controls

4.9.5.1 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.

4.9.5.2 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.

4.9.5.3 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.

4.10 O7A Decoder - Cube

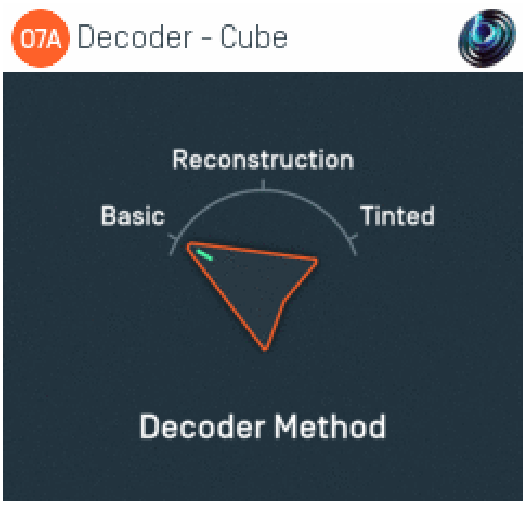

4.10.1 Host Support

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

4.10.2 Audio

	Channels	Content
Input	64	O7A
Output	8	Speaker feeds for cube (mapped to 7.1 for AAX)

4.10.3 Controls

- Method

4.10.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as speaker feeds for eight speakers set out at the corners of an cube.

4.10.4.1 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

4.10.5 Controls

4.10.5.1 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.

4.11 O7A Decoder - Hexagon

4.11.1 Host Support

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

4.11.2 Audio

	Channels	Content
Input	64	O7A
Output	6	Speaker feeds for hexagon (mapped to 6.0 for AAX)

4.11.3 Controls

- Method

4.11.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as speaker feeds for six speakers set out at the corners of an hexagon.

4.11.4.1 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.

4.11.5 Controls

4.11.5.1 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.

4.12 O7A Decoder - Octagon

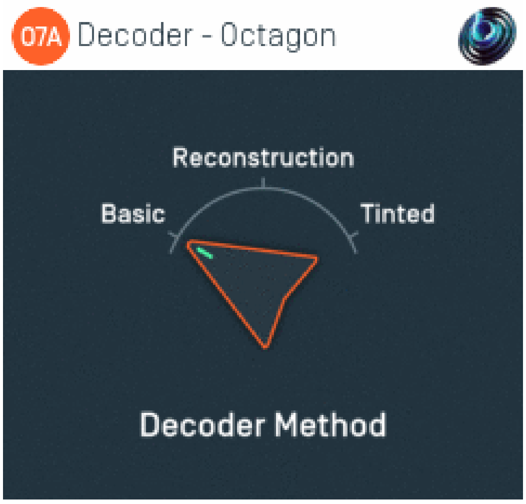

4.12.1 Host Support

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

4.12.2 Audio

	Channels	Content
Input	64	O7A
Output	8	Speaker feeds for octagon (mapped to 7.1 for AAX)

4.12.3 Controls

- Method

4.12.4 Description

This plugin takes a seventh order ambisonics (O7A) mix and decodes it as speaker feeds for eight speakers set out at the corners of an octagon.

4.12.4.1 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.

4.12.5 Controls

4.12.5.1 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.

5 Other Decoders and Converters

5.1 O7A Decoder - Quad Binaural

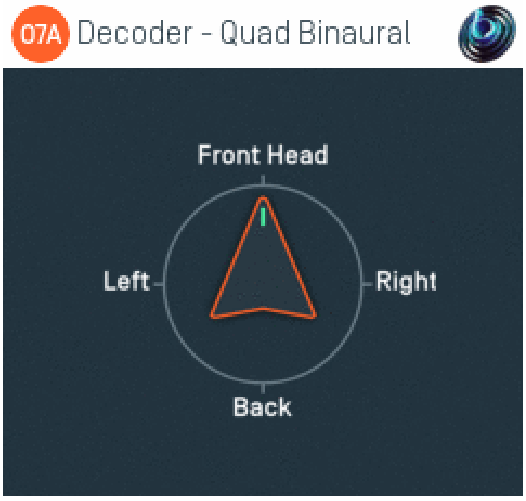

5.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.1.2 Audio

	Channels	Content
Input	64	O7A
Output	2	Binaural Stereo (one head of four)

5.1.3 Controls

- Head

5.1.4 Description

This plugin takes a seventh order ambisonic (O7A) 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 (O7A Decoder - Headphones) instead of this plugin when Quad Binaural is not in use.

5.1.5 Controls

5.1.5.1 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.

5.1.5.2 Technical Notes

This plugin is based on the Blue Ripple Sound Amber HRTF (see the O7A 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.

5.2 O7A Decoder - UHJ Stereo

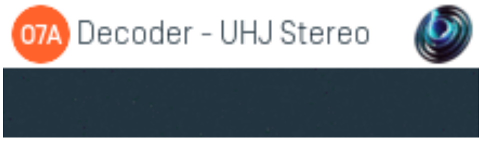

5.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

5.2.2 Audio

	Channels	Content
Input	64	O7A
Output	2	UHJ Stereo

5.2.3 Description

This plugin takes a seventh order ambisonics (O7A) 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 O7A 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 O7A Decoder - Stereo plugin from the O7A Core plugin library. The O7A 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.

5.3 O7A Decoder - N3D

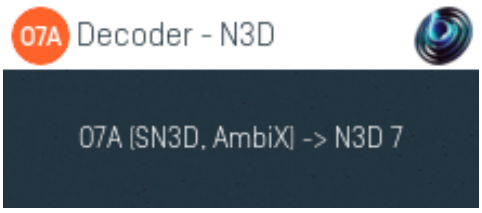

5.3.1 Host Support

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

5.3.2 Audio

	Channels	Content
Input	64	O7A
Output	64	N3D/ACN Ambisonic Audio (mapped to standard SN3D ambisonics for AAX)

5.3.3 Description

This plugin takes a seventh order ambisonics (O7A) mix and converts it to a seventh 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.

N3D is used in MPEG-H 3D Audio.

5.3.3.1 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.

5.3.3.2 Other Conventions

Note that other plugins are available to convert to and from ambisonics using the FuMa convention. See O7A Decoder - FuMa and O7A Injector - FuMa.

5.3.3.3 Orders

This plugin produces seventh order N3D output. If you want lower order output, just take the first channels and ignore the rest. First order needs 4 channels, second 9, third 16, fourth 25, fifth 36 and sixth 49.

Alternatively, the O7A Decoder - O1A, O2A or O3A plugins from the O7A 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.

6 Utilities

6.1 O7A Brickwall

6.1.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.1.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

6.1.3 Controls

- Limit

6.1.4 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 O7A Compressor from the O7A Manipulators plugin library.

6.1.5 Controls

6.1.5.1 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 O7A 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.

6.2 O7A Meter - Karma

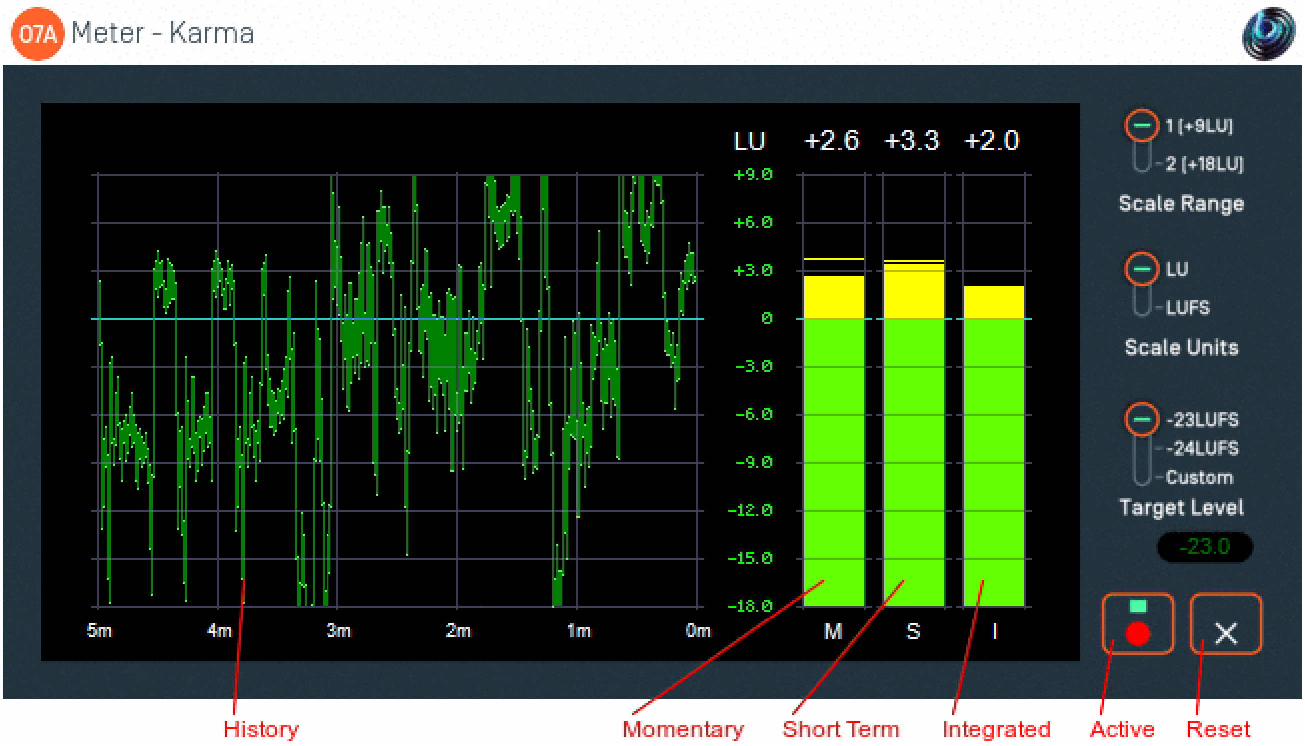

6.2.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.2.2 Audio

	Channels	Content
Input	64	O7A
Output	64	O7A

6.2.3 Controls

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

6.2.4 Description

This plugin measures the loudness of an O7A 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 O7A stream rather than a stream for a specific speaker layout like 5.1.

The O7A decoders included in the O7A Core and O7A 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 O7A mix, it will give you a good idea of the loudness you will achieve later on when you decode for a particular speaker layout.

6.2.4.1 Modifying ITU-R BS.1770 for O7A

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 O7A 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 O7A 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.

6.2.4.2 "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.

6.2.4.3 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.

6.2.5 Controls

6.2.5.1 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.

6.2.5.2 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.

6.2.5.3 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.

6.2.5.4 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.

6.2.5.5 Control: Reset

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

6.2.6 Presets

Presets are available for this plugin:

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

6.3 Stereo Crosstalk Cancellation

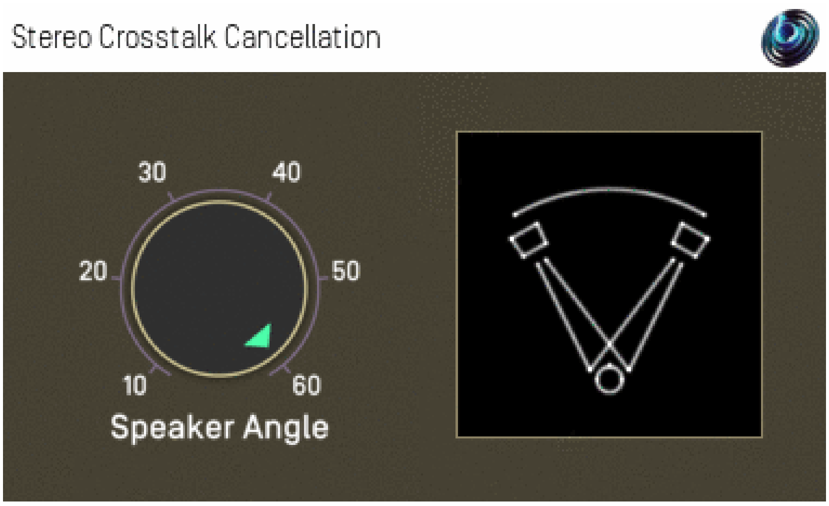

6.3.1 Host Support

Host Type	Support
AAX	Yes
VST2	Yes

6.3.2 Audio

	Channels	Content
Input	2	Stereo
Output	2	Stereo

6.3.3 Controls

- Angle

6.3.4 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).

6.3.4.1 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.

6.3.5 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.

6.3.6 Controls

6.3.6.1 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.

7 Appendix: O7A Streams

7.1 What is an O7A Stream?

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

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

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

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

7.2 What processing can I apply to an O7A stream?

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

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

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

7.3 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.

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

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

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