


The 360pan suite

VST3 - Reaper manual, Mac and Windows



360pan suite is a VST3 and AAX plug-in suite for the Mac designed to deliver ambisonics mixes for immersive audio: audio from all directions, intended for headphone reproduction, that stays where it is even when you turn your head. This manual focuses on the VST3 plug-in for use in Reaper. There is another manual that focuses on the AAX plug-in for use in ProTools HD.

Immersive audio will mostly be required for 360 video projects, and is best enjoyed with VR goggles and or headphones. The ambisonics can be 4 channels (1st order, currently best for YouTube) or 9 channels (second order, recommended for Facebook). Higher orders will be available soon.

The 360pan suite is designed such that you can do most of your work: panning, distancing and mixing right from within the video window in Reaper, where you see the (equirectangular projected) 360 video, without ever needing to open the plug-in interfaces.

The 360pan suite consists of the following plug-ins:

360pan - for the panning and the distancing of sound sources

360monitor- look around and listen to your immersive mix while you are making it

360reverb - convolution reverb with ambisonics impulse responses

360limiter - ambisonics, lookahead peak-limiter with adaptive automatic release time determination, and added EBU R128 (LUFS) loudness metering.

360radar - shows you the location of your audio in your ambisonics mix.

360turner - rotate or tilt an ambisonics recording so that misalignment and calibration errors of the microphone can be easily corrected.

360converter - converts between common ambisonics channel orders and normalisation standards.

This manual

Most parameters of the 360pan suite are self explanatory.

There are a few things that are easy to overlook but are important nonetheless.

This manual first goes into those details, and then discusses each product separately.

The goal of the 360pan suite is to pan source sounds into the ambisonics mix, combine them with real ambisonics recordings, and listen to all this using the 360monitor. In between there is acoustics added, based on distance of the sources, this is provided by the 360reverb. You can verify where in the ambisonics mix your audio is sounding using the 360radar, which displays audio onto the video window.

Example projects and tutorial video

The example sessions, available at audioease.com/360, are 9 channel ambisonics sessions (second order ambisonics). The example session comes in two versions: START and DONE.

The example DONE session does the following:

- it has an equirectangular version of a 360video on a track against which audio should be mixed.
- it has mono dry dialogue that should be dynamically positioned on the people that move around, using automation.

- the dialogue is reverberated dynamically via automation to simulate actors moving away or closer.
- there is a footsteps track that has been recorded with an ambisonics microphone and has been converted to 4 channel b-format, it is piped into the 9 channel bus by simply adding the four channel file to a ten channel output track in Reaper.
- this footsteps track is reverberated in the same space as the dialogue by using a send to the reverb track.
- all audio is combined in the 360monitor that allows you to look and listen around. You can turn by dragging in the 360Monitor video area, or by employing a head tracking device mounted to your headphones (see paragraph headtracker).

Most of this covered in the 10 minute workflow tutorial video:

<https://www.audioease.com/360/#example>

If you want to work along with this video, please open the example START project and follow the steps in the video.

If you have done this once, you have most of what the 360pan suite version 2 can do.

if you want to know the last 30% then read the following in the plug-in descriptions section of this manual.

Essential Reaper settings

When you use the VST3 360pan suite plug-ins in Reaper you will have to verify some settings in Reaper's preferences are properly set up, else the 360pan suite will not function properly.

In Reaper choose Preferences... from the bottom of the Options menu and then go to:

1. Audio -> Buffering -> Anticipative FX Processing should be OFF

2. Audio -> Buffering -> Allow live FX multiprocessing should be OFF

these first two will ensure the 360reverb to function properly (else you will get sync problems)

3. Plug-Ins -> VST

Under VST Compatibility please set the Parameter automation notification to:

Process all notifications

this will ensure automation is written to track when you control the panner pucks from the video overlay (without the panner plug-ins being open).

4. Editing Behavior -> Automation

When recording automation and stopped should be set to: Do not add additional points.

This to ensure you can nudge the time cursor and then drag pucks and write automation this way (like in the tutorial/example project).

5. Reverb last

When creating a bus for 360Reverb, make sure it is last, or right before the monitor.

6. Your Video window can not be 'Docked', and the video window can not be 'full screen'.

These are not in the settings window. You right-click on your video to undock it if it is embedded in another window.

The video window



When you bring up the Reaper video window use the top slider to arrange visibility of items on screen. By color, or per panner puck. Make sure the grid is the size of the video, you can adjust the size of the grid by grabbing the bottom horizontal line of it and dragging it. The top row are individual pucks.

For the overlay to function Reapers Video window must be a separate window, so it should not be 'Docked', and the video window can not be 'full screen'. Right-click on your video to undock it if it is embedded in another window.

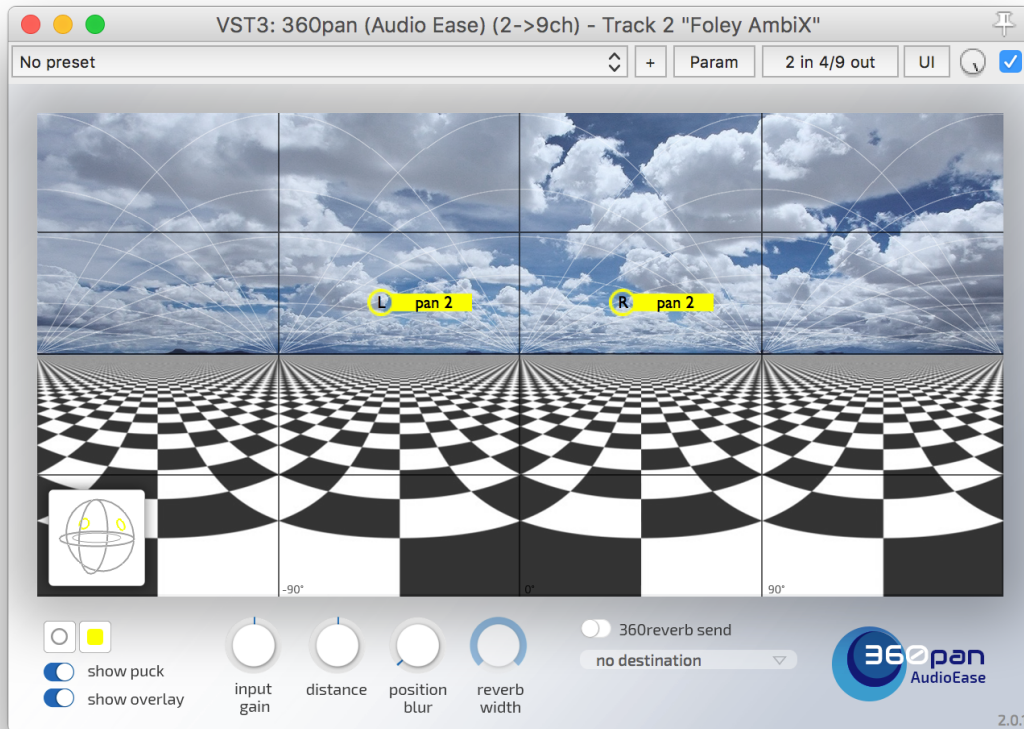
the ® symbol called monitor is a 360radar inserted on aux channel called monitor. The color bar hides or shows pucks with identical colors on or off.

Export

Please refer to <http://www.audioease.com/360> for advice on attaching your audio to video, injecting it with the necessary metadata and uploading it. That information outdates too quickly to be put in this pdf manual.

The plug-ins

360pan




360pan enables you to pan mono, stereo or four channel input files into ambisonics, using pucks that move on top of the Pro Tools video window.

Insert a 360pan and then double click the puck to rename it to something useful.

Gain, distance and reverb width, along with the puck itself will show up in the Pro Tools video window when a panner is inserted and 'show puck' and 'show overlay' are switched on. The three sliders appear when you click the three slider icon that shows when you hover the mouse over a puck in the video window.

You can drag the sliders to a spot where they are not in the way by grabbing them in their white see through area.



Signal Flow

It is possible to use the 360reverb in a traditional way using sends and busses, however it is designed to be used in a more clever way.

mono or stereo sources are converted up to ambisonics, while giving it a left-right, up-down, and distance position in the ambisonics bus. A panner is inserted on each sound source track to do this. So panners get mono or stereo (or even 4channels) of input and output ambisonics.

Each 360reverb creates a private ambisonics input, kind of a backdoor, only accessible by the 360pan instances. Each 360reverb shows up as a send destination in each 360panner. This ties the 360pan's to the reverb in such a way that you can adjust distance and reverb-width per panned source. This goes for an unlimited amount of sources even if there is only a single 360reverb in your project (which is the most usual setup).

Besides all this the usual main input of the reverb still functions (also simultaneously with the private input).

Input gain

The gain slider takes down the audio strength at the input of the panner, so it alters the dry and the reverb of this source, making it the perfect place to mix your audio.

Position blur

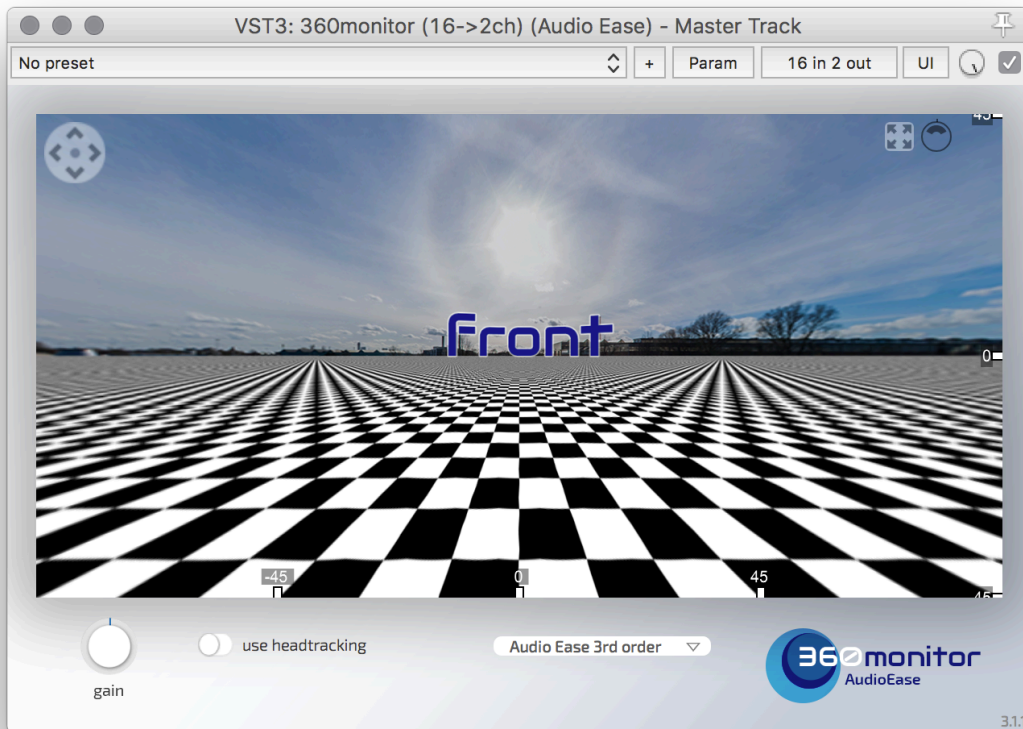
Turn position blur up to create sound-objects that should appear bigger than a point-source. The panning position will then be their center, and they will become larger in the mix.

Head-locking audio

If position blur is turned all the way up, the audio is mono and will not move at all when the listener turns his head. You can use this to mix voice-overs or (mono) music that shouldn't move.

Pan stereo music hard left and hard right and open up position blur halfway to create a head-locked effect in an ambisonics mix, That does not require separate head-locked tracks.

360monitor



360monitor scoops video frames out of the video window and sends them into a browser style drag-n-pan screen while decoding the ambisonics mix to binaural for headphones (or 5.0) for speaker playback, so you can look around and listen to your immersive mix while you are making it.


In the picture the mix format selected is 9 channel (second order ambisonics). This 360monitor is therefore on a ten channel track. It folds the audio down to 2 channels for headphone listening. It does that in the way described technically in Appendix 2 'Ambisonics to binaural conversion'.

Use the expand button in the 360monitor's top right corner to enlarge the screen.

Head tracking

Audio Ease sells a (40 USD) head tracking unit that can be attached to your headphones and can steer the 360monitor. Appendix 3 describes the use of the head tracker.

360reverb



360reverb is the first truly omnidirectional convolution reverb. A single instance can provide pan-following reverb with individually adjustable reverb width for an unlimited amount of sound sources.

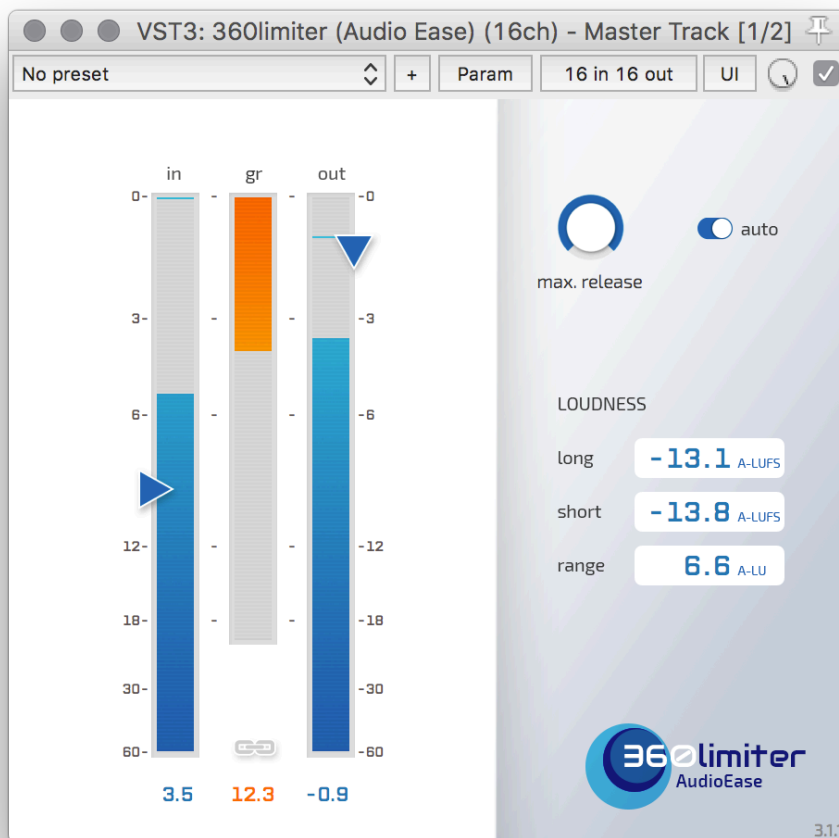
The dry/wet balance and the spread controls only function on the main (traditional plug-in) input of the reverb. 360reverb accepts ambisonics input and delivers ambisonics output. There is a separate private input that is selectable in all 360pan plug-in instances.

The dry/wet balance and the spread controls do nothing on the private audio lines each 360pan has to the reverb. The distance and reverb-width parameters for the private inputs are set in the 360pan plug-in (or through the popup faders in the video overlay).

When creating a bus for 360Reverb, make sure it is last, or right before the monitor.


360limiter

360 limiter is an ambisonics, look ahead, brick wall, peak limiter with adaptive automatic release time determination, and added EBU R128 (LUFS) loudness metering. It will peak-limit first order to third order ambisonics audio in such a way that when gain is reduced, all audio stays in its place. For finding the peak in the ambisonics signal all ambisonics channels are measured. It is a common misconception that one can cut corners by only considering the W-channel: many real life ambisonics signals actually have louder peaks in the other channels.



Live ambisonic recordings often have a large dynamics range, where occasional peaks can prevent the user from gaining without clipping the peaks. When uploaded online (or on any other delivery platform) this can lead to a soft playback. This is where the limiter comes into play. The 360limiter can, transparently, attenuate these peaks, enabling the user to gain the program without clipping.

Operation:

- create an aux with both an ambisonics input and an ambisonics output through which your mix will be routed. Insert a 360limiter in this aux.
- Play your ambisonics through 360limiter while listening to its output. Make sure 'auto' is switched on, and max release is at its maximum. These are the optimal settings to allow 360limiter to adaptively pick the perfect release time.
- Grab the link button  and roll it down until you see the central gain reduction (gr) meter register reduction. Roll down further until you hear unacceptable degradation of your audio material.
- Roll up a little bit, say 3dB, so that your material is safe.
- Now push the out-slider (ceiling) all the way up to get the loudest unclipped and clean ambisonics audio, or less than all the way up in case you aim at a certain LUFS requirement for delivery.

Loudness metering

The Loudness metering is performed on a decoded ambisonics signal. The decoding is performed to a tetrahedron shaped speaker setup called A-format, and the units are therefore called A-LUFS, short for *A-format Loudness Units Full Scale*.

Details about this procedure are in Appendix 4 of this manual.

The loudness information is provided to give an objective indication of the perceived loudness of the material. Various delivery platforms have already set out guide-lines to what is considered an accepted loudness level (usually in long and sometimes short term maximums).

Target: -20 A-LUFS

We confirmed that our own “**long term -20 A-LUFS**” mixes sit very well next to premium 360VR content on YouTube and Facebook, like:


- “Pearl” from the Google spotlight 360 content on YouTube
- “Through the Ages” from National Geographic on Facebook

Controls:

Everything above the threshold control in the in-meter will be limited down to beneath that threshold. Nothing will exceed the threshold. The ceiling control then pulls that threshold level up after limiting. Usually you want the output ceiling at 0dB, and the threshold such that you reach -20 A-LUFS on the long term loudness meter.


Release time is the time it takes for the limiters peak gain reduction to no gain reduction. The release can be set, by the user, from 20 ms to 1000 ms, but can also be set to 'auto', which makes the release dependent on the dynamics of the material it receives. When switching to 'auto' mode the control determines the maximum release time the auto algorithm will choose. Usually 'Auto' is the best choice.

360radar



360radar shows you when there is audio in your ambisonics mix or recording, but more importantly it shows *where* it is, all right in the video window. The 360turner together with 360radar allows you to rotate or tilt an ambisonics recording so that misalignment and calibration errors of the microphone can be easily corrected.

360turner



The 360turner together with 360radar allows you to rotate or tilt an ambisonics recording so that misalignment and calibration errors of the microphone can be easily corrected.

360converter

Converts between common ambisonics channel orders and normalisation standards (FuMa / AmbiX).

Appendix 1

Ambisonics

There are two things you must understand about ambisonics:

- it is the channel layout of your mix, so you mix to ambisonics
- you never listen to ambisonics directly, instead you 'decode' it for a certain type of listening, like headphones (binaural) or four-speakers.

It is not required to understand anything else about ambisonics to be able to use it.

An ambisonics mix describes a full 3d sphere of audio. If in it a sound source appears on a certain position, that position will be more accurate, more precise, when the order of ambisonics is higher. Higher order ambisonics use more channels.


Appendix 2

Ambisonics to binaural conversion and decoder matrices

This final section describes in technical detail the choices that have been made for converting ambisonics down to binaural audio for headphones in the 360monitor. You do not need to understand any of it to use the 360pan suite, this is added for reference purposes.

4 channel to binaural conversion

To give you an idea what you are listening to when you choose 4 channel ambisonics this is it:



The head is you. You are wearing this immensely impractical headphone that consists of a complete rack of speakers. When you turn your head all the speakers turn with it, just like your headphone. looks heavy. And then there is the cowboy that is playing back through that speaker array.

It'd be nice, if you turn your head (and the attached speaker array) the sound world would stay where it is, which as far as the mix is concerned, is identical to panning

the sound world the other way. That requires a mix that can be turned, and tilted, easily. Which is the only reason 360 VR audio often uses ambisonics. it is a mix output format that can be turned easily.

Because the speaker array stays fixed relative to your head, we can replace the whole thing in software by a bunch of impulse responses. All we need is an impulse response from every speaker to every ear and then we fake the whole array of speakers using convolution. Those impulse responses (IRs) are often called HRTFs. So we have a turning mix into a virtual speaker array that delivers to headphones. Whatever way you are looking at it, you are pushing in ambisonics, 4 channels for the simplest version of that, and pulling out headphones, which is 2 channels. It turns out you can meta-sample this 4 in 2 out system, kind of bundle up all these HRTFs inside, so that you end up with 4 times 2 is 8 impulse responses that convert from ambisonics to binaural. We at Audio Ease call those **ambinaurals**, which is the first unusual choice. Others call them “spherical harmonic HRIR filters”.

Ambinaurals are constructed from Head Related Transfer Functions (HRTFs) such that there is one Ambinaural IR for each ambisonics channel per ear. Eight for 4 channel first order. We have used Sadie Binaural Measurements of Google VR SDK so that your 4 channel ambisonics audio will sound identical to the YouTube smartphone app playing back your project.

Second and third order ambisonics to binaural conversion

To give you an idea what you are listening to when you choose second or third order ambisonics check the picture below of the third order decoder. It is a 6-8-6 set of rings. Second order uses a 6-6-6 set of rings with a voice of God.



Bad ambinaurals can make your work as hard as writing a letter with a paint brush

We wanted influence on this most vital stage in the process so we made our own ambinaurals. In the process we wanted to design a system that could capture individual ambinaurals quickly.


Resulting ambinaurals perform well for many listeners, seem to be as neutral as can be expected while still offering dramatic placement, even quite some in the vertical plane. We struck our own balance between neutrality and positioning accuracy.

Ambinaurals are constructed out of HRTFs (Head Related Transfer Functions) that are individually recorded from single speakers in the speaker array to each ear of a dummy- or a real head. We employed many known as well as a few unknown strategies for creating HRTFs and then, where theory and practice clashed, we always picked better-sounding over theoretically-superior. This makes for a few odd choices that are explained here:

All selected ambinaurals are made in the Audio Ease semi-anechoic chamber using the Neumann KU100 dummy head as a microphone, which turned out to work better than any of the other real people and dummy heads we tried.

A diffuse field equalisation curve can be obtained by applying a minimum phase inverse of the average of the power spectra of all HRTFs. However, if that is done, the internal free field equalisation curve built into the ku100 head largely undone. We chose not to apply a free field equalisation, because we really like the Neumann KU100's internal diffuse field EQ.


We used bi-amped non-coaxial high quality studio monitors (Genelec 1031) that are well equalized on axis,



Magnitudes per frequency of an Impulse Response of length 480 samples (10 msec) taken from a Genelec 1031 to a Knowles miniature omnidirectional microphone in the Audio Ease anechoic chamber. It stays within 3 dB throughout the frequency band that ends up in the final IR.


which meant negligible make-up filtering was required. We made rather lengthy HRTF's of 512 samples, which on the one hand gives quite a bit of low frequency resolution, but on the other includes room influences of our anechoic chamber that we could do without. So we did end up synthesizing low frequencies, which is usual, by crossing over to a filtered spike that was phase synced per HRTF, but only below 200 Hz, which is lower than usually employed.

The resulting ambinaurals that we picked are non-symmetric. All tiny measurement errors in the build of our speaker configuration, the KU100 model used, aiming of our turning installation etc are still in there. Using symmetric HRTFs by mirroring the left hemisphere for the right ear is more usual. However we did not prefer this.



Left and right speakers to both ears show the differences that a non-symmetrical set of ambinaurals preserves.

To make the process of sampling a person last less than 10 minutes for 121 speaker positions, we employed a turning stool and 5 speakers, using overlapping exponential sine sweeps that had been enveloped to match the speakers abilities best.



Right hand side: the actual speaker setup that is the result of the audio ease sampling procedure, and left hand side: a single vertical stack artificially extended to place the speakers on a sphere. The speaker setup used in the 360monitor in 9 channel mode has been picked from these positions.

A minimum phase inverse filter of an average of the 5 speakers to a free hanging omni-mic at the head's centre position was applied to take out any of the level and spectral differences between those 5 speakers, and also some of the room characteristics. Although we measured many Headphone Transfer Functions as well, the function of which is to take out the end-user's playback system, none of them were applied to these ambinaurals. In the future we may be supplying these for a number of popular mixing headphones.

To cross over from measurements to synthesized low frequencies Linkwitz Riley filters were employed. The Inter-aural Time Delays (ITD) as well as the Inter-aural Level differences (ILD) are completely preserved, throughout synthesized and measured spectra. To deal with the problem that when you turn your head your ears move out of their positions we cut the start of the HRTFs at first incidence at either left or right ear resulting in what we call a front-split speaker matrix, which inserts a fake little vertical gap between the left and right sides of the speaker matrix that is as large as the width of the head.

When judging headtracked ambisonics to binaural conversion there's a few things to pay extra attention to: It should sound crisp and full bodied wherever the sound source is panned, it should be able to place sounds outside of your head at a distance, and when you turn your head each sound should stay in its place. With this in mind for second order we chose three stacked rings of 6 equally spaced speakers (starting at 30 degree azimuth) and added a voice of God. For third order we chose a horizon of 8 speakers, with a central front speaker, and an upper and lower ring of 6 speakers. This can be seen in the picture on page 12.

Stereo and quad output decoders are created using the Ambisonics Decoder Toolbox by A. Heller, using the rEmax values with speakers at a 90 degree angle from each other on the horizon.

For 5.0 speaker decoding we use the angles from ITU 5.1. We use the HF coefficients of the ITU 5.1 preset from AmbDec by F. Adriaensen as the decoder.

Warning: It is the opinion of Audio Ease that ambisonics the way it is produced for 360video sounds a lot less than ideal on speakers. In fact we think it doesn't sound anywhere near professional on speakers. This is because ambisonics is very fragile: it works in a perfect setup, while sitting at the exact sweetspot. The reverbs we recorded have been made using an ambisonics microphone, which is a coincident system consisting of cardioid microphones, Audio Ease much prefers to use spaced omnidirectional microphones for speaker playback of acoustics, which is the default system for our Altiverb convolution reverb. We recommend against using the output of the '360monitor to speakers' as an end result for anything. We offer it for convenience, to allow you to mix your work while monitoring on speakers rather than on headphones.

Appendix 3

Using the head tracking

A wireless head tracking device can be connected to your headphone to give you the head tracking experience in conjunction with the 360monitor.

This is how it works:

Your mac must be equipped with Bluetooth (or a Bluetooth adapter)

Order your headtracker module from www.audioease.com/store. If you find it elsewhere then do make sure the battery is included, because often it is not. (It should contain the terms MPU6050 or MPU9250, Bluetooth and 9-axis in its description)



Charging the head tracker

Use the included usb cable to charge the internal lithium-ion battery of your head tracker from any usb port, hub or loader. Red light means that it is charging, no red light means device is fully charged

Installing the head tracker




use velcro, a tie-wrap or sticky tape to attach the headtracking device to the top of your headphones, button on the left side.




Switch the device on by sliding the knob to the left (a blinking blue light should appear, if not revert to 'charging the head tracker').

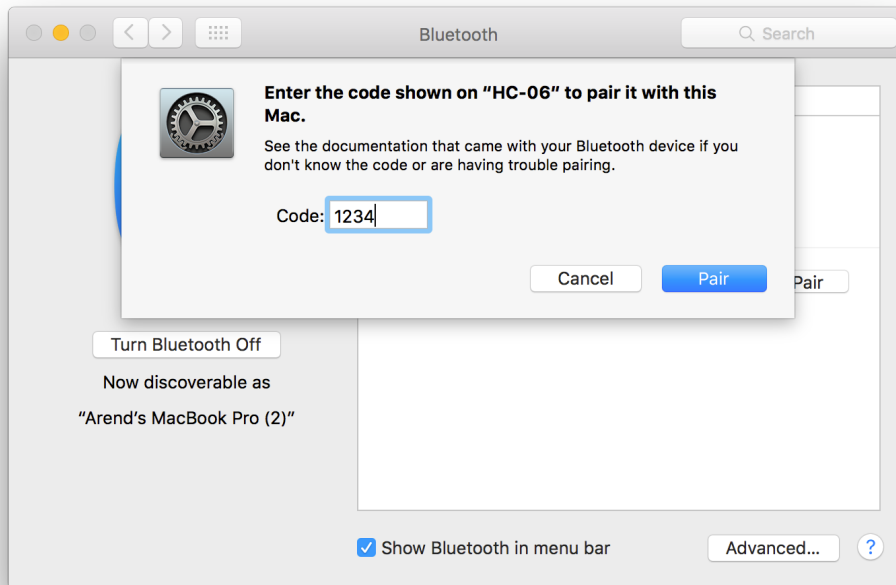
Bring up system preferences->Bluetooth and press Turn Bluetooth On



It takes a little while before HC-06 shows up. Press Pair on the HC-06 item




Press Options




Enter 1234 and press Pair (or Connect)

After a few seconds this is what it then looks like:



Insert a 360monitor and switch on use headtracking
At this point the head tracker should be blinking the blue led.



Press calibrate and follow on screen directions.

At this point the headtracker should be displaying the non-blinking (solid) blue led.

Trouble shooting

- We are getting between 2 and 3 hours of headtracking from a single battery load. At the end of this period the tracking gets sloppy or we experience 'drift' slow turning while we are not turning.

To remedy this just charge your Headtracker.

- If you get drift or sloppy tracking when the unit is fully charged, please re-calibrate by pushing calibrate. This time try to make bigger, slower longer movements covering more of the sphere you wish to be able to reach. The calibration is looking for a center of that sphere.

Leds on the head tracker

There are two leds (lights) next to the on/off slider on the head tracker that show the current status of the device. The device is turned off with the slider in the right position.

- blinking blue: device is on but not connected, possibly paired.
- constant blue: connected and paired (open the 360monitor plug-in and switch on the headtracking option for this).

While device is usb connected:

- red: charging
- no red: device is fully charged

Appendix 4

Loudness Unit Measurement considerations

Loudness Units aim to emulate human loudness perception. A -23 LUFS master should sound as loud as any other.

B-format is not to be played back, it needs to be unpacked to speakers or headphones first. That is called decoding, and Audio Ease selected to decode to speakers for their LUFS measurements.

Any ambisonics decoder has an internal fixed volume knob. We set it to be neutral for uncorrelated sound fields which means that we set our LUFS decoder's volume knob such that when an uncorrelated sound field of -23db LUFS is at the input, and we pan that into B-format and decode it into the simplest speaker configuration (a tetrahedron) it comes out -23db LUFS as well.

There seems to be no disagreement on how loud a sound is supposed to be panned into b-format. More precise: most current b-format panners (or encoders as they are sometimes called), use $W=1$, which corresponds to Ambix SN3D normalisation. This decision leads to a -3dB square wave being panned right in front of you will read -3dB on the W channel as well as on the X channel.

With that we can decide on our volume knob in the decoding step:

We calibrated the decoder volume to $1/4 \sqrt{2}$, or - 9 dB.

To sum it up:

The decoder used for LUFS measurement in 360limiter decodes to four channels on the angles of a tetrahedron (this can be called A-Format). It is the simplest 3d decoder imaginable.

We use these formulas

$$FLU = 1/4 \sqrt{2} (W + Y + Z + X)$$

$$FRD = 1/4 \sqrt{2} (W - Y - Z + X)$$

$$BLD = 1/4 \sqrt{2} (W + Y - Z - X)$$

$$BRU = 1/4 \sqrt{2} (W - Y + Z - X)$$

The above is a valid 4 channel surround format so we can measure LUFS according to the standard ITU BS17703-03*

We call the result A-LUFS.

A-LUFS is the LUFS measured on the A FORMAT signal that is decoded from the Ambix SN3D B-format input of the plug-in, where the decoder gain is set to -9dB.

answers to possible questions:

Why call it A-LUFS, why not just LUFS?

LUFS are described in the standards document ITU BS17703-03. B-Format is not a multichannel signal in the spirit of the ITU BS17703-03 standard document, because the document assumes that each channel of a multichannel signal corresponds to a speaker. B-Format must be decoded before it complies with that requirement. So what we do is decode the b-format to the simplest multichannel speaker format that uses the X Y and Z angles. It is called A-format. We then say: *This B-FORMAT signal file has this amount of LUFS when decoded to A format.* It is vital information that the B-format has been decoded to A-format

before measurement, and therefore we call the unit: A-LUFS, short for A-Format Loudness Units Full Scale.

How do Audio Ease's A-LUFS compare to LUFS reported by Mix Loudness of the Facebook 360 Spatial workstation ?

If a stereo mix that is -10 LUFS is panned hard right and left into ambisonics you can put our meter and Facebook's meter on the ambisonics track and you will see that Audio ease reports -10 A-LUFS, and Facebook reports -7 LUFS. So Facebook reports 3 loudness units (or 3dB) higher than we do. This depends on the exact input material. Facebook Usually reports 3 to 5 db louder than Audio ease. This also means that when Oculus requires -14 LUFS, your Audio Ease target should be -18 A-LUFS.

The viewer looks around which changes the level on his headphones, why don't you measure that ?

- Firstly we would need to decide where the user looks. This cannot be told from the audio. The only thing we can do is find the loudest viewing direction which is a rather random assumption: birds in the trees on a walk in the forest, or audience nearby in a concert setting, a tour guide next to you pointing out buildings in the valley below, a soccer match, none of those requires the user to look into the loudest direction.
- In addition we would need to decide on a maximum human head turning speed.
- Thirdly we would need to know which HRTF's are employed in the decoding process. There are many different ones, Facebook's, Youtubes', personalised ones, and then different ones for different orders of ambisonics. on the fourth place: a user can also look around in a standard 5.0 mix in a theatre as well. The LUFS of the 5.1 mix does not depend on the viewers viewing direction, but only on the audio itself. This idea we extended to ambisonics.
- in the real world perceived loudness (not measured peak !) does not increase when you view in a source's direction, so it shouldn't even matter where you look.
- It would also violate the spirit of the LUFS measurement standard. In a theatre rear audio is to be weighted + 1.5 dB according to LUFS measurement standards. By determining a viewing direction (usually forward) and applying binaural filters you would get an attenuation of the rear rather than an amplification.

How was the decoder and the volume calibration done ?

- We calibrated an uncorrelated sound field by producing 4 separate streams of white noise and panning those at the tetrahedron angles into B-format

These angles are:

noise channel 1 (FLU) azimuth= -45 elevation= +35
 noise channel 2 (FRD) azimuth = 45 elevation = -35
 noise channel 3 (BLD) azimuth = -135 elevation = -35
 noise channel 4 (BRU) azimuth = 135 elevation = +35

We employ industry standard SN3D Ambix b-format panning (encoding) that uses the following formulas to go to B-format.

W = 1;
 Y = sin(azimuth) cos(elevation)
 Z = sin(elevation)
 X = cos(azimuth) cos(elevation)

The above algorithm is used in all of the 4 ambisonics panners we tested, including our own.

The B-format becomes:

$$W = \text{FLU} + \text{FRD} + \text{BLD} + \text{BRU}$$

$$Y = \text{FLU} - \text{FRD} + \text{BLD} - \text{BRU}$$

$$Z = \text{FLU} - \text{FRD} - \text{BLD} + \text{BRU}$$

$$X = \text{FLU} + \text{FRD} - \text{BLD} - \text{BRU}$$

The B-format is then converted to a tetrahedral 4-speaker setup using a gain of -9dB, (which corresponds with a factor of $1/4 \sqrt{2}$)

$$\text{speakerFLU} = 1/4 \sqrt{2} (W + Y + Z + X)$$

$$\text{speakerFRD} = 1/4 \sqrt{2} (W - Y - Z + X)$$

$$\text{speakerBLD} = 1/4 \sqrt{2} (W + Y - Z - X)$$

$$\text{speakerBRU} = 1/4 \sqrt{2} (W - Y + Z - X)$$

The output speaker channels. should then measure the same LUFS as the output noise channels.

Does the limiter respond to LUFS as well ?

No. The goal of the limiter is different: you want to reach a A-LUFS goal while making sure nothing clips in the B-Format master. This may require peak limiting. All channels of the B-format are considered and the entire B-format is limited such that spatial integrity is maintained (everything stays in its place)

* https://www.itu.int/dms_pubrec/itu-r/rec/bs/R-REC-BS.1770-3-201208-S!!PDF-E.pdf