ZEPHYR: AMBISONICS BINAURAL/SURROUND REVERB

1 INTRODUCTION

Zephyr is a 3D audio binaural/surround algorithmic reverb plugin based on higher-order Ambisonics (HOA). Ambisonics is a speaker-independent format capable of representing 3D sound-fields suited for directional or spatial audio. By designing an algorithmic reverb that operates directly within the Ambisonics domain, we extend the immersive quality and control of binaural and surround systems, parameterizing important spatial cues of early reflections, scattering patterns, and reverb distributions. The result is a complete spatial audio package for crafting 3D soundscapes in music, virtual/augmented reality, and beyond.

2 INSTALLATION

To install the Zephyr plugin, unzip and then copy & paste either the 32bit or 64bit VST2/AU files into the plugin directories or paths that your DAW searches for. E.g.

x86 or Win32 build (Windows 7+)
C:\Program Files (x86)\Common Files\VSTPlugins

x64 build (Windows 7+)
C:\Program Files\Common Files\VSTPlugins

mac universal build (i386 + x86_64) (OSX 10.7+)
/Library/Audio/Plug-Ins/VST/

AU build (OSX 10.7+)
/Library/Audio/Plug-Ins/Components/

3 REGISTRATION AND UPDATING

The demo version of the plugin will have periodic silence every 30 seconds without a registration function. The full version can be downloaded after purchase and uses a simple serial-key system. Copy & paste the entire key as provided in your receipt into the supplied key.txt file in the Zephyr folder

e.g. YOURNAME-AAAAA-BBBBB-CCCCC-DDDDD-EEEEE-FFFFF-GGGGG

Then save and close the file. Next, copy/paste the new plugin over any older versions in directories that your DAW searches for. Launch your DAW and load the plugin, click the register button and browse/select for key.txt. If successful, the bottom panel will display your information. Restart or reload your plugin to ensure that the registration was successful.

4 WORKFLOW

The Zephyr plugin integrates three key features from Higher Order Ambisonics (HOA) format with digital signal processing reverberation techniques:

1. **Multichannel mono-source encoder:** Up to 8 sources can be spatially positioned along a Mercator panner and encoded into 1st to 3rd order HOA format per plugin instance. Additional spatial attributes such as distance, source width/size, and strength are parameterized to give fine-grain control over the early spatial impression. Last, a multi-stage chorus processor assigned to each source can further break-up timings to give a sense of movement.
2. **Ambisonics reverb**: A feedback-delay architecture in the Ambisonics domain controls how encoded sound-sources distribute into early reflections, spatially move in time, cohere/dissolve into late reflections, and diffuse into a reverberant tail. Frequency dampening and decay filters are exposed so that material absorption properties can be matched.
3. **Binaural and surround-format decoders:** Multiple Ambisonics decoders are configured to render to headphone (HRTF-based) and surround-speaker formats. An HRTF dataset is bundled within the plugin. An Ambisonics sound-field rotator is provided for efficient alignment adjustments targeted at physical speaker layout.

4.1 MERCATOR SPATIAL PANNER

Zephyr's spatial panner is based on the Mercator map projection (sphere projected onto unrolled cylinder) commonly used in 360 authoring suites. Sound-sources positioned over the 2D Mercator grid are spatially encoded according to their spherical coordinates (azimuth along the horizontal axis, elevation along the vertical axis in Figs. 1 and 2). For the reference binaural decoder, the intersection of the azimuth and elevation axes refer to the direction of the nose; the bottom margin collapses onto the point beneath the head, the top margin onto the point above the head, and the left and right margins onto the meridian running behind the head. The L and R indicators refer to the directions of the two ears.

Sound-sources are positioned via left-click/drag movements by the mouse and will automatically wrap around the grid if the cursor leaves the panner. A selected or active-source is circled and its parameters shown in the source-panel to the right. Scrolling with the mouse-wheel cycles the active-source. Right-clicking a source will toggle rendering. Changes to the active-source's position and rendering state are reflected in the knobs/switch within the source-panel which offer fine-grain tuning and automation.
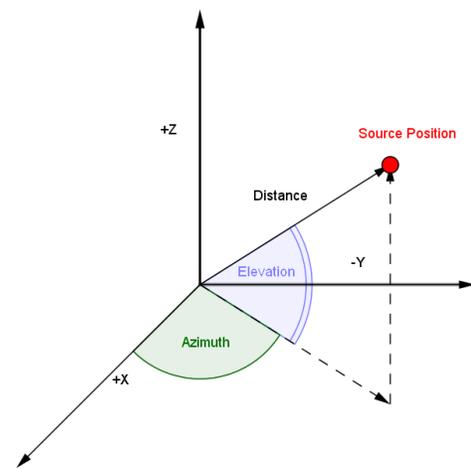

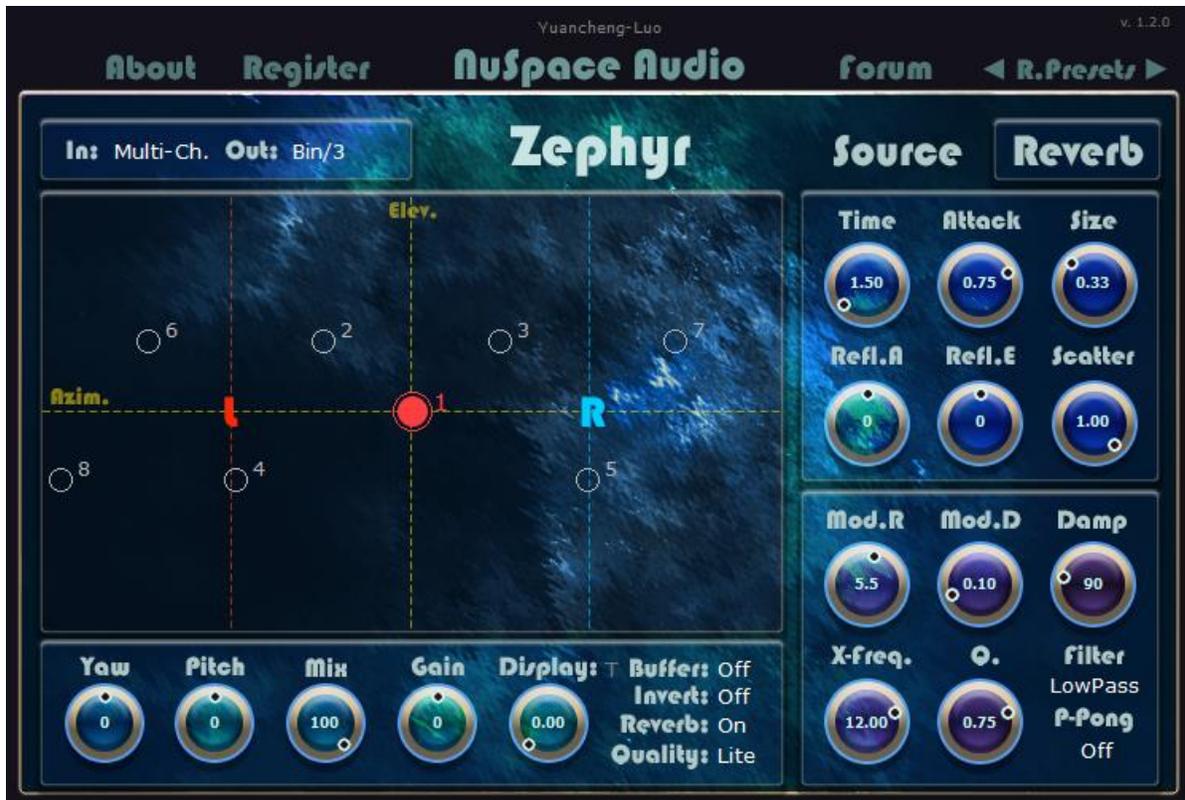*Fig 3: Source position in spherical coordinates*

Fig 4: Colored indicators represent active sound-sources spatialized on the Mercator panner

4.2 GENERAL MOUSE AND KNOB ADJUSTMENTS

Knobs will increment/decrement by vertical mouse drag movements or by scrolling with the mouse-wheel. Holding down shift or control while doing this enables fine-grain adjustment. Double left-click will reset the value to the default. Right click will open a text-box for inputting exact values. Source-direction (azimuth, elevation) and sound-field yaw rotation knobs have the full 360-degree wrap-around range for convenient automation.

4.3 PRESETS

Two independent preset browsers (one for sources, one for reverbs) are hard-coded into upper-right menu; source presets only modify parameters in the source-panel as do reverb presets in the reverb-panel. Clicking on the left/right arrows will scroll through the available presets. Saving presets must be done in the format native to your DAW. Extra or buffer parameters are pre-allocated within the parameter lists to account for both feature additions in future updates and backwards compatibility with older preset configurations saved on file.

4.4   CHANNEL CONFIGURATION

4.4.1   Input: 16 channels

4.4.1.1   Multi-channel: *16 mono streams*

4.4.1.2   Mono Only: *Stream from only channel index 1*

4.4.1.3   Ambisonics (AmbiX): *1ˢᵗ to 3ʳᵈ ACN order-SN3D normalized channels [1-4, 5-9, 10-16] respectively.*

4.4.2   Output: 16 channels

4.4.2.1   Binaural HRTF: *Decodes 1ˢᵗ to 3ʳᵈ order Ambisonics reverb into virtual speakers before rendering over left and right HRTF fields. Outputs to channels 1 and 2.*

4.4.2.2   Stereo/Binaural Hybrid: *Decodes 1ˢᵗ to 3ʳᵈ order Ambisonics reverb into virtual speakers before rendering over left and right interaural time-delay and gain fields (no coloration). Outputs to channels 1 and 2.*

4.4.2.3   Ambisonics (AmbiX): *Bypass decoding and directly output Ambisonics in ACN order-SN3D normalization into channels 1-16. Channels above the Ambisonics order are zeroed.*

4.4.2.4   Stereo (Left, Right): *Left and right speakers are positioned at 9 and 3 o'clock respectively. Outputs to channels 1 and 2.*

4.4.2.5   Quadraphonics Grid (Front left, front right, back left, back right): *Adjacent speakers are separated by 90-degrees intervals. Outputs to channels 1-4.*

4.4.2.6   ITU 5.0 Grid (Front left, front right, front center, surround left, surround right): *No LFE. Outputs to channels 1-5.*

4.4.2.7   FILM 5.1 Grid (Front left, front center, front right, surround left, surround right, LFE): *Silent LFE, Outputs to channels 1-6.*

4.4.2.8   SMPTE 5.1 Grid (Front left, front right, front center, LFE, surround left, surround right): *Silent LFE, Outputs to channels 1-6.*

4.4.2.9   Hexagon Grid (Front left, left, bottom left, bottom right, right, front right): *First speaker positioned at 11 o'clock with subsequent speakers at 60-degree intervals. Outputs to channels 1-6.*

4.4.2.10   Octagon Grid (North west to North East counter-clockwise): *Adjacent speakers separated by 45-degree intervals. Outputs to channels 1-8.*
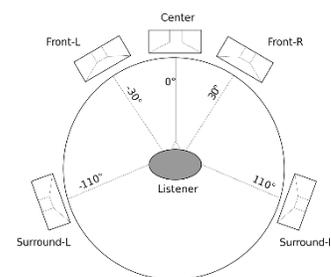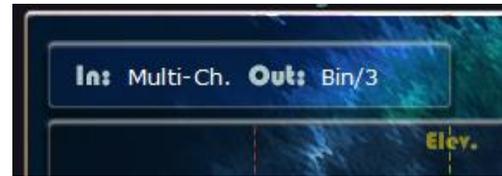
5 CONTROLS

5.1 SOURCE PHYSICS: SIZE, POSITION, AND DELAY MODIFIERS

5.1.1 **Delay:** Lag-time of input stream before direct sound-source and reverb kick in [0 to 1000 ms]

5.1.2 **Strength:** Amount of pre-gain applied to input stream [-infinity to +12 dB].

5.1.3 **Width:** Radius of the sound-source. Larger values turn a point-source's plane-wave into a volumetric source's near-field. [0 to 1] dimensionless.

5.1.4 **Azimuth:** Point-source spherical coordinate component angle from [-180 to 179] degrees.

5.1.5 **Elevation:** Point-source spherical coordinate component angle from [-180 to 179] degrees.

5.1.6 **Depth:** Large values increase the distance between listener and sound-source. [0 to 1] dimensionless.

5.1.7 **Channel Index:** Input stream number [1 to 16] to render. Note: If input format is mono, all sources default to channel 1 regardless of selection.

5.1.8 **Enable Switch:** Toggles rendering to source on and off. Turn off for unused sources to save CPU cycles.

5.2 SOURCE FREQUENCY: CHORUS AND TIMBRE MODIFIERS

5.2.1 **Multi-Chorus Processor:** Pre-processes the input stream before rendering to both the direct sound-source and reverb components.

5.2.1.1 *Chorus Modulation Rate: Frequency of all chorus processors. [0 to 10] Hz*

5.2.1.2 *Chorus Modulation Depth: Larger values increase oscillation depth of all chorus processors. [0 to 1] dimensionless*

5.2.1.3 *Chorus Stages: Modifies the number of chorus processors [0 to 7]. Setting to 0 disables the processor and saves CPU cycles. Larger values add additional copies and increase the signal gain so compensate with the strength control accordingly.*

5.2.2 **Low Cut:** Modifies the resonant-free high-pass filter cutoff frequency [0.06 to 4] kHz.

5.2.3 **High Cut:** Modifies the resonant-free low-pass filter cutoff frequency [0.06 to 22] kHz.

5.3   Reverb Physics: Shape, duration, reflection, diffusion

5.3.1   Time: RT60 duration from [0.1 to 99] seconds of the reverb low-end.

5.3.2   Attack: Modifies the rate that early reflections lose coherence and blend into the late reverb. Smaller values produce discrete echoes, delaying the onset of a diffuse tail. [0 to 1] dimensionless.

5.3.3   Size: Modifies the duration between the onset of early reflections and subsequent gaps in discrete reflections. Larger values will also increase the time for diffusing into a late-tail. [0 to 1] dimensionless.

5.3.4   Reflection Controls
Distribute scatter pattern of early and discrete reflections

5.3.4.1   *Azimuth Angle:* *Modifies the delta change in the azimuth angle component of adjacent early and discrete reflection from [-180 to 179] degrees.*

5.3.4.2   *Elevation Angle:* *Modifies the delta change in the elevation angle component of adjacent early and discrete reflection from [-180 to 179] degrees*

5.3.4.3   *Scatter Rate:* *Large values increase the rate that early and discrete reflections diffuse, forming a smooth late-reverb tail. If both attack and scatter controls are set to 0, discrete reflections retain their coherence and produce an echo train. [0 to 1] dimensionless.*

5.3.4.4   *Ping-Pong Switch:* *Enable to force adjacent discrete reflections to reflect across the origin.*

5.4   Reverb Frequency: Modulation and Material Properties

5.4.1   Modulation Rate: Modifies the modulation rate of reflections from [0 to 10] Hz.

5.4.2   Modulation Depth: Larger values increase the modulation depth. Set to 0 to disable. [0 to 1] dimensionless.

5.4.3   Dampening: Modifies the high-frequency decay curve of air at Nyquist in [0 to 400] dB/sec.

5.4.4   Material Absorption: Additional filter for adjusting the frequency decay curve

5.4.4.1   *Filter Type:* *Low/high pass, low/high shelf, resonant modes*

5.4.4.2   *Cross-over frequency:* *Modifies the central frequency of the material filter from [0.06 to 16] kHz.*

5.4.4.3   *Filter Quality:* *Modifies the material filter's bandwidth. [0 to 1] dimensionless*

5.5 Miscellaneous

5.5.1 Yaw: Rotates the acoustic-field about the Z axis. [-180 to 179] degrees

5.5.2 Pitch: Rotates the acoustic-field about the Y axis. [-180 to 179] degrees. Note that setting both Yaw and Pitch to 0 disables the rotator and saves some CPU cycles.

5.5.3 Mix: Modifies the percent of the wet to dry signal sent to the output. The behavior varies according to the type of input/output configuration.

Configurations (wet/dry)	Out: Headphone/speakers	Out: Ambisonics
In: Multichannel/mono	Decoded reverb + sources / Decoded sources	Amb. reverb + sources / Amb. encoded sources
In: Ambisonics	Decoded reverb / Decoded Amb.	Amb. reverb / Input Amb.

For multichannel/mono inputs, the wet output sums both reverb and source; dry output contains only source. For Ambisonics inputs, the wet output contains only reverb; dry output contains only source.

5.5.4 Output Gain: Modifies the output gain of the wet signal only. [-60 to +60] dB

5.5.5 Display: Modifies the acoustic field display sensitivity. Set to 0 to turn off. [0 to 1] dimensionless.

5.5.6 Background Transparency: Toggle for transparent or black panner background. [T, B]

5.5.7 Buffer Switch: Turn on for fixed block-size processing (1 buffer latency) for smoother source/rotation movements. Turn off for zero latency processing if DAW uses fixed buffer-sizes.

5.5.8 Invert Switch: Turn on to reflect the sound-field about the origin. Use this switch in conjunction with the Yaw rotation (set to 0 or 180 degrees) to realign input sound-fields if necessary.

5.5.9 Reverb Switch: Turn off to render only the direct source / input sound-field and bypass the reverb component (reduce CPU cycles).

5.5.10 Reverb Quality: Set the reverb's quality (density) and how discrete reflection evolve over time:

Quality	Density	Discrete reflections	CPU Load
Lite	Baseline	(A=azimuth, E=elevation)	1x
High	Dense	(A, E)	2x
High Bifurcation	Dense	(A, E), (-A, -E)	2x
Ultra	Smooth	(A, E), (A, E), (A, E), (A, E)	4x
Ultra Clover	Smooth	(A, 0), (-A, 0), (0, E), (0,- E)	4x

5.6 MIDI Values Correspondences (Angular Values)

Parameter Type	MIDI Value	Physical Value (degrees)
Azimuth	0	-180

	MAX	179
	MAX/2	0
Elevation	0	-180
	MAX	179
	MAX/2	0

Canonical planes	Source azimuth MIDI	Source elevation MIDI
Horizontal plane level to nose	Any	MAX/2
Median plane intersecting nose	MAX/2	Any