MUZE: BINAURAL REVERB / 3D PANNER-MIXER

© NuSpace Audio

1 INTRODUCTION

Muze is a multi-reverb VST/AU plugin that combines a 3D audio mixer with a high quality binaural near & far field model. The mixer component enables tight control over the early spatial impression of the acoustic field while psychoacoustic models add essential directional cues to help localize sound. Good externalization and depth is achieved with our reverb design that integrates with binaural cues, giving appropriate contrast for the early spatial impression to sit in.

Muze supports up to 8 input channels with an additional ninth summation channel that doubles as a delay unit. Channels are associated with point sound-sources in an interactive graphical display and have separate processing paths that come together in a shared reverb component. Create whisper effects, wide impressions, impossible spaces. Up-mix mono sources into stereo/binaural. Virtualize 5.1+ to headphones by placing virtual speakers on the sound-stage. Switch to the panning mode for speaker setups.

2 INSTALLATION

To install the Muze 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

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

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

Then save and close the file. Next, launch the plugin, click the register button and browse/select for key.txt. If successful, the bottom left panel will display your information. Restart or reload your plugin in your DAW to ensure that the registration was successful.

4 General Workflow

The Muze plugin routes up to 8 input channels from your digital audio workstation (DAW) into virtual point-sources placed within a spherical sound-stage (see Figure 1). Individual sound-sources can either be cycled to via the mouse-wheel or individually selected via left/right click. Routing is DAW specific whereby mono and stereo channel pairs can directly be mapped to the 8 input channels via inserts or alternatively through side-chaining (see example in Figure 2). Note: Only 2 of 8 channels are set to support side-chaining to avoid overlapping with standard 5.1 surround.

Figure 1: Selecting sound-sources and toggling the mute switch to enable

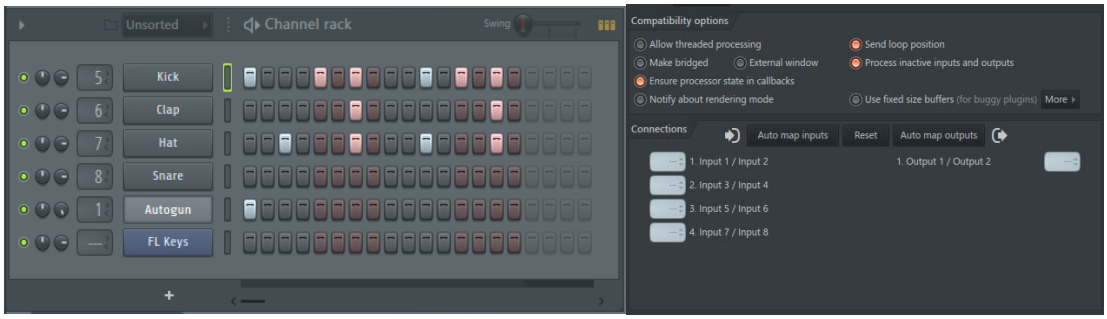

Figure 2: Channel routing from various inserts into the first 8 channels

4.1 Source and Listener Positioning

Sound sources are positioned around a listener centered at origin; the sound-stage displays an overhead view of the listener's head facing the +X axis. Each sound-source's position is specified in terms of spherical coordinates (azimuth, elevation, distance) relative to listener and can be adjusted via click/drag on the sound-stage or via precise knob controls within specific ranges to cover the entire field: Azimuth and elevation [-180 to +180] degrees, distance [0 to 3] meters (See Figure 3).

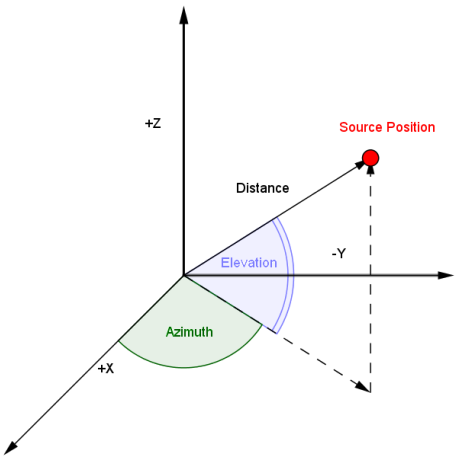

Figure 3: Spherical coordinates specify sound-source placement

The listener's orientation is adjustable for two out of three major principal rotation axes (Yaw and Pitch, no Roll in Figure 4)

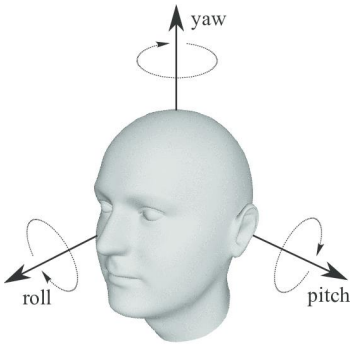

Figure 4: Principal axes orient the listener

For ergonomics and visualization, changing the listener orientation will update the display by rotating all the sound-sources on the sound-stage instead (Figure 5). Last, processing for each channel can be toggled on/off via the mute switch.

*Figure 5: **Top left**: Default positions and orientation with 9 sound-sources circling the listener on the same listening plane; sound-source 1 is directly in from of the listener. **Top right**: The listener turned 90 degrees clock-wise; the left ear now faces sound-source 1. **Bottom left**: The listener pitches backwards where sound-sources previously behind the listener are now higher in perceived elevation. **Bottom right**: Sound-source 1 has been rotated 180 degrees behind the pitched listener.*

4.2 Far & Near Binaural Field Placement

The direct sound-source (without reverb) is binaurally processed into stereo output channels according to the relative position/orientation of the source and listener where separate head-related transfer

function (HRTF) models are applied. The near-field model occurs from 0-1 meters and exaggerates the interaural intensity differences (IID) (especially low frequencies). The far-field ranges from 1-3 meters where a customized HRTF model is used. Note for direct sound-sources, the distance-to-listener effects only the signal gain (no low-pass, no added delay) as to preserve the original signal and avoid inter-source phase misalignment.

4.3 Spatialization Switch

If the spatialization switch is enabled, then direction and frequency-dependent HRTFs will color the sound due to the acoustic wave scattering off parts of the head-torso-pinnae before entering the ears. If the spatialization switch is disabled, then only the average IID between the ears due to head-shadowing are used (left and right signals are only gain-modified). In both cases, the sound-source will be processed with delay according to direction-dependent interaural time differences (ITD) between the ears.

4.4 ITD Scale

The ITD model is designed to fit the average human head whose size can be scaled from 0 (removed entirely or no delay), to 50x (unrealistically large) to achieve interesting time-varying delay effects. More on this with regards to spin and modulation controls in the next sections. For stereo speaker (non-headphone) configurations, it is recommended to set ITD to 0 as to avoid phase-cancellation.

4.5 Hemisphere Switch

The hemisphere switch toggles between two ITD models. When enabled, time-delays of sound-rays are applied only to the contralateral ear which allows for a blending between modulated effects and the original source-source. When disabled, a full spherical delay model is used which will exaggerate phase-change type effects as the sound-source moves or the listener turns.

4.6 Spin and Y-Axis Controls

The spin knob controls the automatic rotation of the listener head, either along the yaw axis (Z-axis default pointing above the listener) or along the pitch axis (Y-axis if the YAx switch is enabled). In the neutral position (0), the listener will orient itself to the specified yaw and pitch coordinates in the GUI. Otherwise, the listener will spin around the axis of choice in terms of revolutions per second (-16 to +16 rev/sec). Returning the knob to 0 (double click) will reorient the listener back to the neutral position.

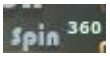

The range of rotation (from 0 to 360 degrees) is specified by the number to the right of the spin label which can be adjusted by dragging up/down or by right-click and entering an exact figure. For values less than 360, a sin LFO drives the rotation about the specified yaw or pitch coordinates by +/- value degrees. For value equals to 360, the rotation remains unidirectional as if the listener never stopped turning.

Dialing-in different spin rates can be used to achieve effects from smooth rotational panning of the entire sound-field around the listener, swish-effects of fast moving objects, and sponge/flange/modulation effects by adjusting the ITD and reverb panning settings. Note: Only the sound is updated here (sound-stage visualization remains unchanged).

4.7 HRTF MODULATION SWITCH

The modulation setting effects whether HRTF rendering is applied either pre-or-post reverb. Normally, HRTFs applied at the ends of the signal processing chain produce a psychoacoustically accurate image (modulation switch off) but we can invert this relation to achieve interesting effects within time-varying systems. For example, enabling the modulation switch with non-neutral head spin, pan, and ITD can realize pitch-shifting, warble, flange-like effects within the reverb component. **Note:** If the spatialization switch is disabled and HRTF modulation enabled, then a quasi-HRTF model is used which produces a more immersive spatial image than having disabled both HRTF and spatialization switches.

Spatialization (Column) **HRTF Mod (Row)**	On	Off
On	**HRTF → Reverb:** Experimental effects, 100+% performance	**Reverb → Quasi-HRTF:** Panning with strong coloration of early spatial image, 80% performance
Off	**Reverb → HRTF (default):** Most immersive, 100% performance	**Reverb Only:** Panning with transparent reverb, 50% performance

Table 1: HRTF and reverb rendering order and performance

4.8 THE 9TH CHANNEL

The 9th sound-source is a feedback delay unit that processes the sum of all non-muted input channels (1-8). Its output delay (1 to 16 taps per second) is displayed to the upper-right of the "dist" label

and can be adjusted by vertical mouse-drag. The frequency decay rates are identical to the reverb unit specified by the damp and RT60 knobs. Azimuth, elevation, and distance functionalities remain identical to the other channels and each instance of the delay effect feeds into the reverb unit. Use the delay unit to create echo-type effects that can fade into the distance and combine with spin controls to achieve 3D ping-pong sounds.

5 CONTROLS

5.1 GENERAL MOUSE AND KNOB ADJUSTMENTS

Knobs will increment/decrement by **vertical** drag movements of the mouse 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 defaults. Right click will open a text-box for inputting exact values. Positional and orientation knobs (azimuth, elevation, yaw, pitch) have the full 360-degree range with wrap-around for convenient automation.

5.2 CHANNEL SELECTION IN THE SOUND-STAGE

Scrolling the mouse wheel on the spherical sound-stage will cycle through the nine channels/sound-sources. Clicking on individual sound-sources will set that sound-source to the active status. Only the active channel's 3 positional parameters (Azimuth, Elevation, and Distance) and its mute status are shown and updated in the GUI. Right-click on an active sound-source will toggle the mute status.

Dragging a sound-source will translate its position along the listener oriented XY plane; for non-zero elevation, the source position will fall towards the center-of-head as you move it towards the center of the sound-stage. For automation, positional parameters for sound-sources can be independently linked.

6 REVERB SETTINGS

The reverb component models the acoustic properties of various spaces depending on the selected algorithm or mode. In general, the dimensions of the space are scaled to be within 5-23 meters, adjustable along each X-Y-Z axes, and centered around the listener at the origin. The space's orientation is given by the absolute X-Y-Z frame of reference as shown in the sound-stage when the listener is in its default orientation. Some reverb controls are shared across all algorithms, some unique.

6.1 REVERB MODES/ALGORITHMS

6.1.1 Velvet: Room/Hall Reverb

Velvet independently models specular and diffuse type reflections in a room-like enclosure. The relative contribution of each reflection type is adjustable as both add to the early & late components of the reverb. Specular reflections impact both ears (mono) whereas diffuse reflections are decorrelated to produce a wide impression. The room size effects the onset delay and the perceived density of the late reverb. A small room with an unnaturally high RT60 time will sound overly dense whereas a large room will have a sparser / bouncy feel to it (low echo density). Use the seed knob to tweak or randomize resonances due to the formation of room-modes.

6.1.1.1 Specular Reflections

Specular reflections are mirror-like or ray-traced acoustic events of the pressure wave scattering off the flat boundaries of a room, constituting both early and late reflections. The specular control knob modifies the base decay rates of various reflection-trains computed from the RT60 reverb time and the room size (0 returns the fastest decay, 1 returns the base decay rate).

6.1.1.2 Diffuse Reflections

Diffuse reflections are Lambertian-like or stochastic acoustic events of the pressure wave scattering off rough boundaries of the room and accounting for most of the late-reverberation. The diffuse control knob modifies the base decay rates of the diffusion process computed from the RT60 reverb time and the room size (0 is fastest decay, 1 returns the base decay rate).

6.1.1.3 Quality

Adjusts for CPU performance versus reverb saturation (0 shuts off reverb, 1 gives the lowest quality, 5 the highest). Mid-end systems should be able to handle settings of 2-3 with minimal overhead. High to ultra-quality settings (4-5) are designed for reverb-dominant effects (e.g. long reverb tails, spin/mod effects) and final-renders. Note: Enabling spatialization requires 2x the

computational load; enabling non-neutral spin with spatialization requires 3x the computational load. For solo use of the mixer/panner (no reverb), set the quality to 0.

6.1.2 Lime: Early/Echo Reverb

Lime jointly models early and diffuse type reflections within a single network by controlling both the separation (in both space/time) of early reflections and the rate of cross-mixing that form the late reverb. Well-separated early-reflections can be individually spatialized and extended into long echo trains that may never diffuse (works best with impulsive/transient sounds). Slap-back, trickle, and ping-pong effects are possible. Separate patterns of early reflections are selectable.

6.1.2.1 Discrete Amount

Controls the separation of discrete early reflections in each pattern. A minimum separation of 0 exhibits few early reflections whereas a maximum spread of 1 can convert early reflections into late echoes.

6.1.2.2 Blend Rate

Controls the rate that discrete reflections mix with each other. Large values will fully blend the early reflections into a diffuse tail. Small values may give rise to distinct echoes in the late-tail.

6.1.2.3 Discrete-Pattern Selection

Specifies various configurations of the early reflections in the physical space that depend on the room shape and size. Modes 0-5 correspond with square, rectangle, thin-rectangle, cube, prism, tall-prism lattices respectively.

6.1.3 Twister: Modulation/Wind Reverb

Twister models amplitude and delay-based modulation effects of diffuse reflections within a single network. Amplitude modulation contributes to “whoosh” style effects whereas delay-line modulation produces characteristic “chorus/buzz/swarm” effects. The combination of the two allows various sound-sources to evolve into different wind-noise patterns.

6.1.3.1 Delay Modulation Rate

Controls the modulation rate (0 to 5 Hz) of several delay-lines within the network.

6.1.3.2 Delay Modulation Depth

Controls the number of samples traversed by the delay-line modulators (0 = disable). The slope of frequency shift over time is proportional to the product of the rate and depth with higher-frequencies decaying faster.

6.1.3.3 Wind Rate

Controls the modulation rate (0 to 5 Hz) of several evolving amplitude envelopes that effects the rate of diffusion within the network. Larger values create dual side-bands that extend in +- frequency directions from your input sound-source and grow wider with longer RT60.

6.1.4 Indigo: Shimmer/Frequency-Echo Reverb

Indigo integrates both a pitch-shifter and frequency-shifter module within a feedback loop to produce cascading harmonic/inharmonic sounds. The pitch-shifter preserves the input harmonics by transposing the signal up/down in integer multiples of *K* semitones. The frequency-shifter distorts the input harmonics by moving the signal frequency up/down in integer multiples *L* Hz. An additional echo knob adjusts the flow of the feedback loop, interpolating between shimmer and delay effects suited for soft and/or percussive type inputs.

6.1.4.1 Pitch Shift

Controls the number of semitones that the input signal is transposed every step (-12 to +12) or +- 1 octave. A shift-factor of *K* semitones transposes a signal by multiples of $2^{K/12}$ in frequency.

6.1.4.2 Frequency Shift

Controls the frequency (-1K to 1K Hz) to shift the input signal every step. A factor of *L* carries a signal by multiples of *L* Hz. Use in conjunction with the pitch-shifter to produce a matrix of inharmonics.

6.1.4.3 Echo

Controls the diffusion of reflections within the pitch shifting/sliding network (0 = seamless or diffuse, 1 = discrete or echo). The room-size determines the time-difference between steps.

6.1.5 Solarium: Swarm/Swell Reverb

Solarium produces broadband delay-based vibrato effects within a cascading build-up of diffuse reflections. Vibrato components simulate doppler shifting that spread out in frequency over time to create characteristic swarm/buzz effects. The reverb buildup generates a swell effect suited for ambient pads and choruses.

6.1.5.1 Vibrato Rate

Controls the number of oscillations (0-10 Hz) per second.

6.1.5.2 Vibrato Depth

Controls the number of samples traversed by the delay-line modulators (0 = disable). The slope of frequency shift over time is proportional to the product of the rate and depth.

6.1.5.3 Swell

Controls the mixture between onset reverb and the buildup of cascading reverb (0 for onset only, 1 for slow buildup then decay). The room-size determines the duration of the buildup whereas the decay is controlled by the RT60.

6.1.6 Osmose: Detuned/Water Reverb

Osmose simulates bassy under-water effects capable of producing detuned reverb tails. Unlike modulated diffuse reflections that weaken as they move away from a center frequency, osmosis applies a selective pressure that causes center frequencies to diverge, concentrating the energy instead towards growing positive and negative side-bands. Use this reverb for creating tension and eerie ambient pads.

6.1.6.1 Osmose Rate

Controls the number of delay-line oscillations (0-10 Hz) per second.

6.1.6.2 Opacity (Transparency)

Controls the degree of coloration between the direct and reverb signals. Low values (0) preserves the direct signal whereas large values (1) adds more glue or connectivity in frequency to the early reflections.

6.1.6.3 Osmosis

Controls the bias towards centralized frequency diffusion (0) versus side-banded frequency osmosis (1) during modulation.

6.1.7 Hedera: Dispersion/Lush Reverb

Hedera creates blur-type and dense reflections approaching the noise-band. The blur strength links with the attenuation of the direct-signal by the mask knob, softening any transients in the source. Several modulators link to the twirl control and the reverb tone's high-end can be rolled-off via the Lush knob. Use this reverb for simulating dense natural environments.

6.1.7.1 Twirl Rate

Controls the delay-line modulation rate (0 to 10) Hz/sec.

6.1.7.2 Lush (Tone and dispersion control)

Controls the roll-off high-frequency gain of several filters and the diffusion pattern. Set to 0 for no roll-off with full-diffusion in frequency. Set to 1 with maximum roll-off (-36 dB/sec) and dispersion in frequency.

6.1.7.3 Mask

Controls the blur strength and the attenuation of the direct signal; a value of 1 maximally blurs the early signal whilst attenuating the direct-signal by 90%.

6.1.8 Zyphire: Bloom/Air Reverb

Zyphire fashions bloom (short-swells) and airy-wash effects for more distant sound sources. The bloom effect produces a natural rising crescendo when set high. When set low, early reflections are both spatialized and randomized. Extremes of the frequency spectra can be rolled off using the air knob. Use this reverb for simulating marble and/or hard surface environments.

6.1.8.1 Zyphire Modulation Rate

Controls the delay-line modulation rate (0 to 10) Hz/sec.

6.1.8.2 Air (Low and high-frequency roll-off)

Controls the roll-off rate for both sub-bass and high-frequency decay of airy-washes. Large values maximally lets through sub-bass and high-frequency washes whereas low values aggressively cuts them out.

6.1.8.3 Bloom

Controls the swell strength and the attenuation of the direct signal; a value of 1 maximally crescendos the early signal whilst attenuating the direct-signal by 90%. A value of zero will discretize the early echoes for larger room sizes.

6.1.9 Neptune: Nonlinear/Early Reverb

Neptune customizes the distribution of early reflections to achieve gated and occluded reverb effects. Parameterizing the duration or spread of early reflections without the roll-off of a late reverb produces punchier transients. Adjusting the density of early reflections and diffusion rate of late reflections muffles the direct source to resemble an occluded space.

6.1.9.1 Neptune Modulation Rate

Controls the delay-line modulation rate (0 to 10) Hz/sec.

6.1.9.2 Density (Early reflections and late diffuse)

Larger values increase the number of early reflections, increase the diffusion of the late reverb, and slightly attenuates the direct source.

6.1.9.3 Spread (Early reflections)

Large values space out early reflections in time independent of RT60. For gated reverb effects, reduce the reverb time to 0 and increase the spread to the desired fraction of one second.

6.2 Reverb Time

The overall reverb time is given by the bottom-most right knob in terms of RT60 (0.1 to 99 seconds). Together with the dampening parameter, these controls adjust for the perceived density of reflections according to the room size. With the spatialization switch enabled, the various reflections are processed through the HRTF models in batches, intended to achieve externalization. Note for fine-grain adjustment, hold down shift or control while adjusting the knob or use right-click.

6.3 Dampening

Allows for high-frequencies to decay faster as par normal atmosphere attenuation than the specified RT60 in terms of dB/sec at the Nyquist frequency.

6.4 Pan

This setting allows for subtle modification to the reverb direction. While “directionless” reverb is ideal for contrasting directional cues against the direct sound-source, biases in the room shape, size, and reflectivity properties create perceptible differences as the listener orientation changes. When spatialization is enabled, the pan setting distorts the reflection model along the y-axis giving the impression of more reflective walls and a general bias of late-reverberation in that direction. When spatialization is disabled and modulation enabled, linear-panning laws are used with regards to orientation of the y-axis with respect to the left and right ears. When both spatialization and modulation are disabled, the default y-axis orientation is treated as an absolute axis and will not change respect to listener orientation and spin settings.

6.5 Seed

The seed control in the far upper-right alters the randomization used by different algorithms. Adjust this value to produce different variations in the reverb tail or to circumvent undesirable resonances if your sound-sources are narrow-band or harmonically sparse.

6.6 Reverb Solo

Will output only the reverb component, ignoring the direct sound-sources.

6.7 SENDS

This knob controls the fraction of signal to send to the reverb component (0 is no signal, 1 is full signal). Use it to attenuate the reverb in the case that the reflections drown out the direct sound-source, interfering with its externalization.

7 MISC.

7.1 0-LATENCY

By default, the plugin processes with a 1-buffer latency as specified by the settings in your DAW. However, 0-latency is possible (toggle the switch) with the risk of sound tearing during fast rotations if the DAW does not send fixed-length buffer frames.

7.2 MIX

Mixes the dry signals from only the first two input channels with the two wet output channels. 0 percent outputs only the dry signal only whereas 100 percent produces the wet signal only.

7.3 MIC LOCK

Mix Lock: Click the U/L to the right of the mix knob to lock the mix-ratio ratio while changing presets. “L” = locked, “U” = unlocked.

7.4 GAIN

Controls the final gain (-60 to +20 dB) applied to the wet signal only.

7.5 PRESETS

Factory-presets can be navigated within the GUI or via the DAW’s interface. Modifications to factory-presets selected via the GUI do not override the factory presets themselves. Selections done via the DAW’s interface however make lasting changes until you reset the plugin.

8 DAW SPECIFICS

8.1 FLSTUDIO

Tested with 4 pairs of stereo channel.

8.2 REAPER

Tested with 8 pins. For AU version, Reaper defaults to using the two aux channels (maps to sources 3-4) but can be changed back to 8 pin setup by adjusting the +/- pin buttons.

8.3 ABLETON LIVE

Tested with 4 pairs of stereo channels. Channels 3-4 are specified as aux.

8.4 LOGIC X

Supports only stereo input + single side-chain. Side-chains default to channels 7-8.

8.5 MAX/MSP

For Mac, only 64-bit mode will display the native GUI.

8.6 MISC.

Stereo configurations working on Cubase 5-9 elements, Bitwig, Studio One 3, MuLabs, Cantabile.