# VOLTAVERB

## DSP Overview

*Architecture and design philosophy behind the sound.*

Version 1.0.0

Volta Plugins Ltd

www.voltaplugins.com

# Contents

# 1. Design Philosophy

Voltaverb is built around a central idea: **convolution reverb should be a creative instrument, not just a room simulator**. Every design decision flows from this — from the 4-slot morphing architecture to the way effects interact with the impulse response rather than simply processing the output.

Three engineering priorities shaped the DSP architecture:

### 1. No Compromises on Audio Quality

All internal processing uses 32-bit floating point. Filters use topology-preserving transforms derived from analogue prototypes. Parameter changes are smoothed per-sample to prevent zipper noise. Where pitch shifting is involved, multi-point interpolation preserves harmonic integrity.

### 2. Responsive at All Times

IR-modifying parameters (room size, decay, damping) require background reprocessing of the impulse response. Rather than making the user wait, Voltaverb uses a **real-time bridge** architecture: instant perceptual feedback via real-time filters while the IR catches up in the background. The result is a plugin that feels immediate even when the underlying computation takes time.

### 3. IR-Aware Processing

Many of Voltaverb's processors analyse the loaded impulse response and adapt their behaviour accordingly. Diffusion adjusts its density based on how sparse or dense the IR is. The synthetic brightness engine matches the loaded IR's decay characteristics. Damping applies corrections relative to the IR's natural frequency response rather than as an absolute filter. This means every effect sounds musical regardless of which IR is loaded.

> **Why It Matters:** *This IR-aware approach is why Voltaverb can sound natural across wildly different impulse responses — from a tiny cupboard to a cathedral — without the user needing to re-dial every parameter.*

# 2. Convolution Engine

## Non-Uniform Partitioned Convolution

Voltaverb uses non-uniform partitioned convolution, which splits the impulse response into progressively larger FFT blocks. Small blocks at the head provide low latency; larger blocks for the tail provide CPU efficiency. This achieves significantly better CPU performance than uniform partitioning while keeping latency well within acceptable bounds for reverb applications.

## Crossfading on IR Changes

When an IR-modifying parameter changes — room size, decay, damping, early stretch, brightness — Voltaverb re-bakes the impulse response on a background thread and swaps it into the live convolution with a short equal-power crossfade. For most material, the change is smooth. Because convolution rebuilds the tail from the new IR, a tail that is mid-decay at the instant of the swap may dip briefly before the new sound takes over — this is inherent to convolution reverb, not a defect. The XY morph pad is the exception: it blends between already-loaded IRs in real time with no re-bake. A fully double-buffered convolution that removes the dip entirely is planned for a future version.

## Background Processing

All IR processing — resampling, envelope shaping, normalisation — happens on a dedicated background thread. The audio thread never blocks on IR computation. A timer polls for pending changes and dispatches background work, with debouncing to coalesce rapid parameter adjustments into a single reprocessing pass.

> **In Practice:** *This is why you can sweep the Room Size dial smoothly without audio glitches — the audio thread always has a valid IR to work with, and the background thread prepares the next one.*

# 3. XY Morph System

## Parallel Convolution with Corner Skipping

Four IR slots occupy the corners of a 2D pad. Each loaded slot runs its own non-uniform partitioned convolution, and the outputs are mixed according to the current XY weights. To keep this efficient, any slot whose morph weight is effectively zero — for example, slots at unused corners while the cursor is elsewhere — is skipped entirely, since it would add only silence. A single IR parked in a corner, therefore, costs about one convolution, while a full centre blend engages all four.

## Perceptual Weighting

The raw XY position is passed through a perceptual curve before calculating slot weights. This ensures that movement near the centre of the pad feels proportional and musical — without the curve, the blending would bunch unnaturally at the corners.

Bilinear interpolation then produces four weights, normalised to sum to 1.0 across loaded slots only. Empty slots are excluded from the calculation, so two IRs in opposite corners produce a clean crossfade without dead zones.

## Smoothing

XY position changes are smoothed to prevent clicks during rapid movement. Pad movement triggers a fast crossfade to the newly morphed IR, giving stutter-free response even during aggressive automation.

> **Creative Use:** *Automate the X and Y parameters in your DAW for evolving reverb textures. Corner positions are the most CPU-efficient — a single active convolution — while centre blends engage all four slots, so keep that in mind in CPU-heavy sessions.*

# 4. Room Shaping

## Room Size (Interpolation-Based Resampling)

Room Size uses **interpolator-based time-stretching**, not a phase vocoder. This changes both the timing and pitch of the impulse response — physically accurate to how a larger or smaller space would sound. Doubling the room size lowers pitch by an octave and doubles reflection spacing, just as physics would dictate.

Two quality tiers operate depending on context:

- **Preview mode** (while the dial is moving): A fast polynomial interpolator provides responsive feedback
- **Final quality** (after the dial stops): A high-quality windowed sinc interpolator provides pristine results

A perceptual power curve maps the dial to the stretch factor, providing finer resolution in the musically useful mid-range and broader strokes at the extremes.

| Dial Position | Perceived Effect |
|---|---|
| 50% | Small booth — tight, intimate |
| 100% | Original IR — natural character |
| 150% | Large hall — spacious, open |
| 200% | Cathedral — vast, deep |
| 300% | Massive space — epic, cinematic |

When compressing (stretch < 1.0), a steep zero-phase anti-aliasing filter prevents foldback artifacts. After resampling, energy-based normalisation ensures consistent output level regardless of room size.

Frequently-used size values are pre-cached for instant response. Moving between cache points uses interpolation rather than resampling from scratch.

## Early Stretch

Unlike Room Size, Early Stretch reshapes just the early reflections — the first portion of the IR — for independent control over the room's initial character. It works in the time domain: the distinct early reflections are detected as discrete peaks and repositioned to new arrival times (their position scaled by the stretch amount), while the diffuse energy between them is time-scaled with the same high-quality windowed-sinc interpolator used by Room Size.

Because the reflections are repositioned rather than smeared, their attack character is preserved in a way naive stretching would not manage. Only the early section (before the early/late boundary) is affected; the late tail is left untouched and crossfaded back in, and the stretched early section is energy-matched so its level stays consistent.

> **Key Difference:** *Room Size changes both time and pitch (physically accurate). Early Stretch mainly reshapes the timing and density of the early reflections, keeping them close to the original pitch (a creative tool). Use Room Size for realistic room scaling; use Early Stretch for independent early-reflection shaping.*

## Room Bright

When Room Size stretches beyond 110%, the resampling naturally shifts high frequencies downward. Room Bright compensates by generating new synthetic high-frequency content from the existing IR using harmonic excitation. The exciter blends in progressively — subtle at small stretches, prominent at extreme sizes — ensuring stretched rooms retain realistic air and presence.

## Room Tone

Room Tone adds adaptive low-frequency warmth. It uses FFT analysis of the loaded IR to identify the room's natural bass character, then applies a tailored bass boost that complements the space. IRs with weak bass receive more enhancement; bass-heavy IRs are boosted gently. The effect scales with Room Size, and turning the Room Tone amount up gradually layers in subtle harmonic saturation for additional warmth.

# 5. Decay & Auto-Gain

## Decay Envelope

The Decay control shapes the reverb tail by applying an amplitude envelope to the impulse response. The auto-gain stage measures the decay-shaped IR with a perceptual (K-weighted) loudness model and compensates the level, so perceived loudness stays consistent across decay settings. The first portion of early reflections is always preserved, so Decay shapes the tail without reducing the room's initial character.

Pre-cached decay levels enable instant response. Intermediate values interpolate between cache points in logarithmic space, and real-time smoothing prevents zipper noise during automation.

## Auto-Gain

After all IR processing (morphing, room size, decay, damping, early stretch), the auto-gain system normalises output to a consistent target level. The measurement uses weighted RMS across all active morph slots, with the effective IR length capped to prevent extreme settings from skewing the result.

> **In Practice:** *With Auto-Gain enabled, you can reshape the reverb's frequency balance using the damping controls without unexpected volume jumps — the auto-gain stage normalises level after all IR processing. (Factory presets ship with Auto-Gain off so manual level-balancing stays predictable.)*

# 6. Pre-Delay

Traditional pre-delay implementations either crossfade between two delay lines (causing audible comb filtering) or click during transitions. Voltaverb solves this with a gain-ducking state machine:

1. When the user changes the delay time, the output gain smoothly descends to silence.
2. At the floor, the delay time switches instantly (inaudible because the output is silent).
3. After a brief hold, the gain smoothly rises back to unity.

The hold period coalesces rapid dial changes into a single duck event. If the user continues turning during the hold, the timer resets without re-ducking — eliminating the grating stutter that occurs with repeated rapid transitions.

The delay line uses high-order polynomial interpolation for sub-sample accuracy, supports tempo sync across standard note divisions (1/1 through 1/32 including triplet and dotted values), and includes a fine-tune offset for micro-timing adjustment in sync mode.

> **Why This Matters:** *This duck-hold-rise approach means you can sweep the pre-delay dial freely without hearing comb filtering or clicks — something most convolution reverbs struggle with.*

# 7. Damping

## Differential Design

Voltaverb's damping system applies the difference between the loaded IR's inherent frequency response and the target damping — not an absolute filter. When an IR already contains accurate damping characteristics, the real-time filters return to unity gain, preventing double-processing.

This differential approach means damping always sounds musical regardless of the source IR. You're adjusting relative to the space, not applying a fixed filter on top.

## Multi-Band Architecture

A 3-band damping design (low, mid, high) using shelving and bell filters — with crossover points aligned to the offline spectral-damping bands — gives independent control over each frequency range. Per-band smoothing keeps parameter changes click-free.

## Distance Simulation

The Dist (Distance) control simulates how sound changes as it travels through air using an FFT-based spectral processor applied directly to the IR. Negative values apply progressive, time-increasing high-frequency absorption — later reflections lose more brightness than early ones, just as in a real large space. Positive values add presence and proximity. The frequency weighting curve targets the mid-high range rather than just the extreme top end, matching the physics of real air absorption.

# 8. Parametric EQ

Voltaverb's EQ uses **Topology-Preserving Transform State Variable Filters (TPT SVF)** — the same filter topology found in leading professional EQ plugins. TPT filters are derived from analogue prototypes using the trapezoidal integration method, providing:

- Exact analogue frequency response matching — no frequency warping artifacts at high frequencies
- Inherent stability at all frequencies and sample rates
- Zipper-free modulation with per-sample coefficient interpolation

Available in 4-band (quick adjustments) and 8-band (detailed sculpting) modes, with ten filter types including Bell, Shelving, Baxandall (vintage character), Low/High Pass with selectable slopes up to 48dB/octave, Notch, Bandpass, and Tilt.

The EQ processes only the wet signal, allowing you to shape the reverb's frequency content independently of the dry signal. The interactive frequency display shows the spectrum in real-time with draggable band nodes and an optional phase response curve.

# 9. Dynamics

## Compressor

A full-featured dynamics processor with dual-mode detection (crossfadeable RMS/Peak blend), soft-knee implementation for transparent compression onset, and optional lookahead for transient-accurate gain reduction.

A sidechain high-pass filter prevents bass frequencies from triggering excessive gain reduction — essential when compressing reverb tails that may have significant low-frequency energy.

## Rhythmic Gate

A tempo-synced gate that chops the reverb into rhythmic pulses. Phase-locked to your DAW's transport position across standard note divisions, with adjustable swing for groove, and smooth exponential attack/release envelopes. Depth controls how much the gate attenuates — from subtle breathing to full silence on closed steps.

> **Signal Flow Note:** *The gate is applied after the reverb and effects chain, so you're gating the fully processed reverb tail. This creates cleaner rhythmic textures than gating the input before reverb.*

# 10. Creative Effects

Voltaverb's effects chain processes the reverb output with a suite of creative tools. Each effect is designed to work specifically with reverb material rather than being a generic processor.

## Shimmer

Granular pitch shifting creates multiple simultaneous harmonic layers above and below the reverb signal. Overlapping windowed grains with high-quality interpolation and per-grain micro-detuning prevent the metallic, static quality of exact integer pitch ratios. An internal diffusion network and dynamics processing keep the shimmer smooth and controlled even at high feedback settings.

## Diffusion

A Schroeder allpass cascade smooths out the reverb tail. The delay times are deliberately static, so adjusting Diffusion never pitch-bends the tail. The processor is IR-adaptive: sparse IRs with distinct early reflections receive more diffusion to fill gaps, while already dense IRs receive less diffusion to avoid over-smearing. Prime-number delay times prevent harmonic resonances from building up, and the density comes from the feedback structure and blend curve rather than modulation.

## Freeze

A feedback delay network captures a moment of the reverb tail and sustains it infinitely. The FDN uses an orthogonal feedback matrix for maximum diffusion density, with decorrelated modulation on each delay line to prevent the frozen texture from sounding static or metallic. Stereo-aware routing preserves the spatial image of the captured moment.

## Thickness

Even-harmonic saturation inspired by vacuum tube circuits adds warmth and body. Oversampled processing prevents aliasing artifacts. A built-in mid-band focus EQ concentrates the saturation on the body of the sound rather than harshness in the highs, and integrated compression prevents the saturation from causing excessive level increases.

## Motion

A modulated delay with multi-tap echo creates depth and spatial interest. Multiple independent LFOs at different rates produce complex, non-repeating movement. The processor is IR-adaptive: LFO rates scale inversely with the IR's reverb time (long tails get slower modulation), and echo levels scale with the IR's diffuseness.

## Granular

A full-featured grain engine with up to 128 simultaneous grains, an 8-second circular capture buffer, and extensive per-grain randomisation (pitch, position, direction, filter, timing). Multiple window shapes, variable spray count, and feedback recirculation transform the reverb into evolving, cloud-like textures.

## Vintage (Tape / Vinyl / Tube)

Three analogue-emulation processors add period character to the reverb. Tape applies oversampled tanh-style saturation with low-frequency warmth, gentle high-frequency roll-off, and wow and flutter pitch modulation. Vinyl adds surface crackle, a wear and degradation stage, and low-frequency turntable rumble. Tube provides asymmetric-bias even-harmonic saturation with DC blocking and makeup gain. Each sub-effect has its own enable, and a global Vintage Mix sets the overall dry/wet blend.

## Synthetic Brightness (Air)

An algorithmic reverb network generates synthetic bright content that is blended with the convolution output. The synthetic engine analyses each loaded IR — its early reflection pattern, decay shape, and density — and matches these characteristics so the added brightness blends naturally. Each morph slot maintains its own analysis data, with the highest-weight slot dominating the behaviour as you move across the XY pad.

# 11. Spatial Tracker (Pan Follow)

The Spatial Tracker analyses the stereo input to determine the perceived source position, then positions the reverb accordingly. This creates a more realistic sense of space where the reverb appears to emanate from where the source is panned.

## Dual-Band Tracking

A crossover filter splits the input into low and high bands, each tracked independently with separate weighting. This allows, for example, tracking the high-frequency transients of a vocal while ignoring low-frequency bleed from a kick drum — solving a common problem with naive pan-following implementations.

## Detection Modes

Three detection algorithms suit different material:

- **RMS:** Smooth, averaged detection — best for sustained material like pads and vocals
- **Peak:** Transient-responsive — best for percussive sources
- **Envelope:** Shaped follower with separate attack and release — best for dynamic sources

## Tracking Behaviour

Hold time maintains the reverb position after the source stops, preventing jittery returns to centre. Hysteresis around the detection threshold prevents low-level noise from causing position flutter. Adjustable attack and release control how quickly the reverb follows movement versus how quickly it returns to centre.

Spatial controls include Reverb Focus (tight follow vs diffuse spread), Pan Range (maximum deflection), Centre Bias (gravitational pull toward centre), Source Width (stereo sensitivity), and Reverb Spread (output stereo width).

# 12. Modulation Engine

## Dual LFOs

Two independent LFOs, each with five waveforms (Sine, Triangle, Square, Saw, Random/Sample-and-Hold). Rates run from sub-Hz to 20 Hz. The two LFOs have fixed routings: LFO 1 modulates output level (tremolo), and LFO 2 modulates brightness (filter sweep). Waveform switching uses a brief crossfade ramp to prevent clicks.

## Envelope Follower

Tracks the input signal level and uses it to modulate reverb parameters. The detection operates in the squared domain for CPU efficiency, with user-adjustable attack and release times.

## Modulation Targets

The envelope follower has a Target selector with five settings — Off (the default, no modulation) plus four destinations: Mix (wet/dry balance), Filter (brightness), Time (decay/size), and Gain (output level). The two LFOs use the fixed tremolo and brightness routings described above; a freely assignable LFO routing matrix is planned for a future update.

> **Creative Use:** *Use the envelope follower targeting Mix to automatically duck reverb when the input is loud, maintaining clarity. Set fast attack and medium release for sidechain-style pumping effects.*

# 13. Signal Flow

The complete signal path through Voltaverb:

```
Input L/R
 |
 v
[Pre-Delay]
 |
 v
[Convolution Engine] --- 4 parallel convolutions, one
 |                       per loaded slot, mixed by XY
 |                       weights (zero-weight slots
 |                       skipped). Room size, decay,
 |                       damping, early stretch and
 |                       brightness baked into each IR.
 v
[Real-Time Damping] --- multi-band correction
 |
 v
[Parametric EQ] --- 4 or 8-band TPT SVF
 |
 v
[Effects Chain]
 |-- Shimmer     |-- Diffusion
 |-- Motion      |-- Thickness
 |-- Vintage     |-- Granular
 |-- Freeze
 v
[Dynamics]
 |-- Compressor  |-- Rhythmic Gate
 v
[Spatial Processing]
 |-- Pan Follow  |-- Stage  |-- Width
 v
[Auto-Gain]
 |
 v
[Dry/Wet Mix]
 |
 v
[Colour] --- IR convolution on dry signal
 |
 v
Output L/R
```

Key points about the signal flow:

- **IR-baked processing:** Room size, decay, damping, early stretch, and brightness are all applied to the IR before convolution — not as post-processing effects.
- **Effects chain:** Creative effects process the convolution output, not the input or the IR itself.
- **Colour is last:** The Colour processor applies IR tonal character to the *dry* signal, creating cohesion between the dry sound and the reverb.
- **Auto-gain position:** Level normalisation happens after all processing but before the dry/wet mix, ensuring consistent output regardless of parameter settings.

# Technical Specifications

| Specification | Value |
|---|---|
| Plugin formats | AU, VST3 |
| Internal processing | 32-bit floating point |
| Supported sample rates | 44.1 kHz to 192 kHz |
| Convolution | Non-uniform partitioned |
| EQ topology | TPT State Variable Filter |
| Crossovers | Linkwitz-Riley aligned |
| IR formats | WAV, AIFF, FLAC, OGG, VIR (encrypted factory) |
| Parameters | All fully automatable |
| Latency | Automatic DAW delay compensation |