MUSHIN

Dual Filter LFO Matrix Synthesizer & Sound Processor

Model: MOD-02
Version 1.0.0
C++/Web Hybrid DSP Core

Translating Market Physics Directly into Digital Signal Processing

Mushin Audio
May 2026

Contents

Chapter 1

Introduction & Core Philosophy

Welcome to **Mushin**, a state-of-the-art hybrid C++/Web sound processor that bridges the volatile mathematics of financial markets with high-performance digital signal processing (DSP).

In standard synthesizers, parameters are modulated using isolated LFOs or standard envelope generators. In contrast, **Mushin maps market physics and trading concepts directly to signal processing mathematics**. The audio waveform fed into the plugin represents the market's data stream, and the processing modules act as momentum constraints, kinetic chokes, and liquidity decays.

1.1 The Market-to-DSP Conceptual Map

Mushin's core architecture is built upon four direct translations of market concepts to signal processing:

Information Saturation (S-Curve Saturation)
As a market trend extends, information saturates, buying or selling pressure plateaus, and volume flattens out. In DSP, this is represented by a soft-clipping waveshaper using a hyperbolic tangent ($\tanh$) transfer function:

$$y = \tanh(x \cdot \text{Drive})$$

As simulated trend momentum extends, the `Drive` increases, rounding off the peaks of your audio waveform and narrowing its dynamic range.

Trend Exhaustion (Hard Clipping)
When statistical exhaustion is triggered (representing the absolute end of buying capacity), all remaining dynamic range instantly evaporates. The S-Curve waveshaper bypasses the soft-clipping $\tanh$ transfer function and switches to a hard-clip threshold:

$$y = \begin{cases} \text{Threshold} & \text{if } x > \text{Threshold} \\ -\text{Threshold} & \text{if } x < -\text{Threshold} \\ x & \text{otherwise} \end{cases}$$

This squares off the waveform, generating harsh odd-order harmonics. It represents the mathematical state of "zero remaining liquidity."

Kinetic Choke (State-Variable Filter Modulation)
Tracking the relationship between Potential Energy and Kinetic Energy, the dual filters act as dynamic chokes. As momentum slows down, the filter's cutoff frequency drops, choking off high-frequency energy. During trend exhaustion, resonance (Q) is multiplied exponentially to create a resonant "scream" before the frequency is clamped down.

Information Decay (Quantization Error / Bitcrusher)

When statistical confidence drains, the resolution of information decays. This is modeled by real-time quantization noise and downsampling. At the peak of exhaustion, audio quality disintegrates under its own weight as bit-depth drops dynamically.

Chapter 2

Technical Architecture & Specifications

Mushin is engineered with an advanced hybrid framework combining high-performance audio processing with a highly interactive, responsive web frontend.

2.1 Core System Features

- **DSP Engine:** Written in highly optimized C++ utilizing the **JUCE 8.0.4** framework and native JUCE DSP modules.
- **User Interface:** A premium web UI built with modern HTML5, CSS3, and ES6 JavaScript, rendered via Microsoft Edge/WebView2 on Windows.
- **Asset Management:** All web assets (HTML, CSS, JS) are bundled via CMake's `juce_add_binary_data` into the plugin binary, ensuring zero-latency asset loading.
- **Event Bridge & Visualizer:** Pushes real-time 30Hz high-frequency waveform updates from C++ directly to the WebView canvas via the `emitEventIfBrowserIsVisible` event bridge, which auto-throttles when the UI is closed to save CPU.
- **Build Environment:** Structured for CMake with the MSBuild (Visual Studio) generator. The active development binary and installer output is compiled in the `build2` directory.

Chapter 3

Module & Parameter Guide

Mushin features 8 high-performance DSP modules, fully automatable and exposed to the host DAW via the JUCE `AudioProcessorValueTreeState` (APVTS).

3.1 Module A: S-Curve Waveshaper (Saturation)

The primary saturation stage, implementing information saturation.

Table 3.1: S-Curve Waveshaper Parameters

ID	Name	Range	Default	Description
`gain`	Gain	0.0 to 2.0	1.0	Post-waveshaper output make-up gain.
`drive`	Drive	1.0 to 20.0	1.0	Amplifies input signal into waveshaper.
`exhaustion`	Exhaustion	OFF, ON	OFF	Bypasses soft $\tanh$ for hard clipping.
`threshold`	Threshold	0.0 to 1.0	1.0	Amplitude ceiling for hard-clipping.
`mix`	Dry/Wet Mix	0.0 to 1.0	1.0	Blends processed and dry signals.

3.2 Module B: Dual Filter System

Two independent filters (A & B) routed in Series or Parallel.

- **Serial:** Audio $\rightarrow$ Saturation $\rightarrow$ Filter A $\rightarrow$ Filter B $\rightarrow$ Outputs.
- **Parallel:** Audio splits equally to Filter A and Filter B. Outputs are summed $(0.5 \times A + 0.5 \times B)$.

3.2.1 Filter Models

1. **Clean (TPT):** Topology Preserving Transform filter. Ultra-transparent.
2. **Vintage (Ladder):** Classic 4-pole transistor ladder model. Warm and driven.
3. **Acid (Diode):** 3-pole diode ladder filter with squelchy, raw resonance behavior.
4. **Digital (State Variable):** Crisp, versatile filter for high-precision cuts.

3.3 Module C: LFO & Modulation Matrix

Two sample-accurate LFOs (0.1Hz - 20Hz) offering Sine, Triangle, Saw, Square, and Random (Sample & Hold) waveforms modulate six key filter targets.

Table 3.2: Filter Parameters (Identical for A and B)

ID (Filter A / B)	Name	Range	Default	Description
`filter_[a/b]_cutoff`	Cutoff Freq	20Hz - 20kHz	20,000Hz	Logarithmic frequency cutoff.
`filter_[a/b]_resonance`	Resonance	0.0 to 1.0	0.1	Cutoff frequency resonance boost.
`filter_[a/b]_drive`	Filter Drive	1.0 to 10.0	1.0	Pre-filter input saturation.
`filter_[a/b]_grit`	Filter Grit	0.0 to 1.0	0.0	Post-filter clipping fuzz.
`filter_[a/b]_type`	Filter Type	Choice	Clean	Filter algorithm selection.
`filter_[a/b]_mode`	Filter Mode	Choice	Lowpass	LP, HP, BP, or Notch slopes.
`routing`	Routing	Choice	Serial	Serial or Parallel configuration.

Table 3.3: Modulation Targets (Bipolar range: -1.0 to +1.0)

Target ID	Target Parameter	Modulation Depth / Range
Target 1	Filter A Cutoff	Modulates frequency by up to **+/- 5 octaves**.
Target 2	Filter A Resonance	Modulates resonance by up to **+/- 1.0**.
Target 3	Filter A Grit	Modulates post-filter grit by up to **+/- 1.0**.
Target 4	Filter B Cutoff	Modulates frequency by up to **+/- 5 octaves**.
Target 5	Filter B Resonance	Modulates resonance by up to **+/- 1.0**.
Target 6	Filter B Grit	Modulates post-filter grit by up to **+/- 1.0**.

3.4 Module D: Sidechain Processor (Envelope Follower)

Tracks the amplitude envelope of an incoming source to dynamically duck or pump Mushin parameters.

Table 3.4: Sidechain Parameters

ID	Name	Range	Default	Description
`sc_active`	SC Active	OFF, ON	OFF	Bypasses or enables sidechain.
`sc_source`	SC Source	Choice	Internal	Main feed, DAW Aux, or test tone.
`sc_threshold`	SC Threshold	-60dB to 0dB	-24dB	Amplitude envelope action threshold.
`sc_attack`	SC Attack	0.1ms - 500ms	10.0ms	Envelope detector attack speed.
`sc_release`	SC Release	1ms - 2000ms	100.0ms	Envelope detector release speed.
`sc_amount`	SC Amount	-100% to 100%	0.0%	Ducking (-) or pumping (+) depth.
`sc_target`	SC Target	Choice	Gain	Target parameter: Drive, Cutoff, or Gain.
`sc_hp_freq`	SC HPF	20Hz - 2kHz	20Hz	Detector high-pass rumble filter.
`sc_lp_freq`	SC LPF	500Hz - 20kHz	20kHz	Detector low-pass kick isolator.

3.5 Module E: Noise Generator / Auxiliary Oscillator

Injects synthetic noise textures or pure synth tones.

3.6 Module F: Trance Gate

Chops the output amplitude based on host tempo and 8 selectable rhythmic step sequences (Straight 16th, Offbeat 16th, Classic 1, Classic 2, Four-On-Floor, Galop, Space Gate, Euclidean 5).

Table 3.5: Noise Generator Parameters

ID	Name	Range	Default	Description
`noise_active`	Active	OFF, ON	OFF	Bypasses or enables generator.
`noise_type`	Type	Choice	White Noise	White, Pink, Sine, Tri, Saw, Square.
`noise_freq`	Freq	20Hz - 20kHz	440Hz	Tone frequency for standard oscillators.
`noise_level`	Level	0.0 to 1.0	0.1	Output level.
`noise_routing`	Routing	Choice	Pre-Dist	Inject Pre-Dist, Pre-Filter, Post-Filter.
`noise_fm_mod`	Filter FM Mod	0.0 to 1.0	0.0	Modulate Cutoffs directly using generator.

Table 3.6: Trance Gate Parameters

ID	Name	Range	Default	Description
`tg_active`	Active	OFF, ON	OFF	Bypasses or enables Trance Gate.
`tg_mix`	Mix	0.0 to 1.0	1.0	Dry/Wet ratio.
`tg_rate`	Rate	1/16, 1/8, 1/4	1/16	Step speed relative to host tempo.
`tg_start`	Start	0% to 100%	5%	Attack fade-in ramp per active step.
`tg_hold`	Hold	10% to 100%	50%	Open gate duration per step.
`tg_end`	End	0% to 200%	10%	Release fade-out ramp per step.
`tg_depth`	Depth	0% to 100%	100%	Closed-step attenuation amount.

3.7 Module G: Info Decay (Quantization Error)

Simulates information decay by executing dynamic bit-crushing and sample-rate reduction.

Table 3.7: Info Decay Parameters

ID	Name	Range	Default	Description
`qe_active`	Active	OFF, ON	OFF	Bypasses or enables Info Decay.
`qe_depth`	Resolution	2.0 - 16.0 Bits	16.0 Bits	Continuous amplitude bit-depth reducer.
`qe_downsample`	Time Res	1.0 to 32.0	1.0	Sample-rate divider (Downsampler).
`qe_mix`	Mix	0.0 to 1.0	1.0	Dry/Wet ratio.
`qe_link`	Exhaustion Link	OFF, ON	OFF	Plunges bit depth to 4.0 during exhaustion.

3.8 Module H: Delay Processor (Stereo / Ping-Pong)

High-quality stereo delay line optimized for rhythmic echoes and physical-modeling comb filtering.

Table 3.8: Delay Parameters

ID	Name	Range	Default	Description
`delay_active`	Active	OFF, ON	OFF	Bypasses or enables Delay.
`delay_time`	Delay Time	1ms - 2000ms	300ms	Delay time (skewed first 30% covers 1-17ms).
`delay_feedback`	Feedback	0.0 to 0.95	0.3	Feedback amount.
`delay_mix`	Mix	0.0 to 1.0	0.3	Dry/Wet ratio.
`delay_pingpong`	Ping-Pong	OFF, ON	OFF	Bounces delay left/right in stereo.
`delay_sync`	Sync	OFF, ON	OFF	Synchronizes delay time to host tempo.
`delay_tempo`	Tempo Sync	Choice	1/4	1/16, 1/8, 1/4, 1/2 note subdivisions.

Chapter 4

UI Operation, Presets & Aesthetics

Mushin's interface is structured around state-of-the-art aesthetics and intuitive parameter tracking.

4.1 The Dynamic Theme System

Visual themes are served dynamically by the C++ engine to the web frontend via a virtual `skin.css` sheet. Theme configurations are persisted in `settings.ini` inside your AppData application directory.

1. **Industrial (Default):** Electric Blue (`#00bfff`), Brushed Metal, Orange accents.
2. **Synthwave:** Hot Pink (`#ff00aa`), Neon Cyan (`#00ffff`), Dark Violet background.
3. **Acid:** Poison Green (`#bfff00`), Lime-glow (`#39ff14`), Toxic Dark background.
4. **Firepits:** Lava Red-Orange (`#ff3c00`), Ember Gold (`#ff6600`), Obsidian background.
5. **Ocean Deep:** Deep Blue (`#060b1a`), Tropical Cyan (`#00f3ff`) accents.
6. **Ice World:** Frosted Grey-Blue (`#0f151c`), Glowing Ice Blue (`#66e0ff`).
7. **Dark Hellish:** absolute Black (`#0e0606`), Blood Crimson (`#ff2200`).

4.2 Preset Browser & Search

The header panel manages XML presets loaded directly from disk:

- **Windows Path:** `%AppData%\Mushin\Presets\`
- **macOS Path:** `~/Library/Application Support/Mushin/Presets/`
- **Controls:** Use the dropdown or search field to filter presets instantly. Pressing `Enter` triggers loading.

Chapter 5

Sound Design Recipes

Here are some creative setups to jump-start sound processing inside Mushin.

5.1 Recipe 1: The Acid Flashback

Goal: Resonant squelching acid filter sweeps.

```
# S-Curve Waveshaper
drive = 8.0
exhaustion = OFF
mix = 0.7

# Dual Filters (Serial)
Filter A: LP, Acid, Cutoff = 400Hz, Res = 0.85, Drive = 3.0, Grit = 0.4
Filter B: LP, Clean, Cutoff = 1.2kHz, Res = 0.2

# LFO 1
waveform = Saw
frequency = 2.5Hz
Mod Matrix LFO 1 -> Target 1 (Filter A Cutoff) = +0.7
```

5.2 Recipe 2: The Vocaloid A-O

Goal: Dual sweeping vowel formants mimicking human speech.

```
# S-Curve Waveshaper
drive = 2.0
mix = 0.3

# Dual Filters (Parallel)
Filter A: BP, Vintage, Cutoff = 800Hz, Res = 0.9
Filter B: BP, Vintage, Cutoff = 1800Hz, Res = 0.9

# Modulation Matrix
LFO 1: Sine, 0.4Hz
LFO 1 -> Target 1 (Filter A Cutoff) = +0.5
LFO 1 -> Target 4 (Filter B Cutoff) = -0.5
```

5.3 Recipe 3: Warm Tape Saturation

Goal: Adds analogue-like warm glue.

```
# S-Curve Waveshaper
drive = 3.5
exhaustion = OFF
mix = 0.4

# Dual Filters (Serial)
Filter A: LP, Clean, Cutoff = 8.5kHz (roll-off high harshness), Res =
    0.0
Filter B: Bypassed (Cutoff = 20kHz, Res = 0.0)

# Sidechain Envelope Ducking
sc_active = ON
sc_target = Gain
sc_source = Internal
sc_amount = -15%
sc_attack = 20ms
sc_release = 150ms
```

Chapter 6

Future Roadmap

Mushin's ongoing feature pipeline includes:

6.0.1 Air Gap Suspension (Granular Looper)

Implements an existential pause. When trend exhaustion triggers, the write pointer of a circular buffer freezes while the read pointer loops continuously with an applied Hanning window to prevent clicks, creating an ambient frozen stutter.

6.0.2 Final Stage Limiter / LMC

A dedicated safety brickwall peak limiter / character compressor. Features a `Clean` transparent safety mode (utilizing a 2.0ms lookahead circular buffer) and an aggressive, highly colored `LMC` (Listen Mic Compressor) feedback mode for vintage NYC parallel smashing.

6.0.3 LFO Phase Invert

Adds an invert toggle to the LFO blocks to reverse the sweep polarity ($\phi = -\phi$), enabling mirror-image modulations using a single shared LFO.

6.0.4 Vocal Tract Formant Filters

Physical acoustic formant filter models added directly to the filter type selection, enabling single-control sweeping vowel morphing.