

Reverie 1.13

2026-06-08 · Parallel Minds · Lyon, France

The Native Release

Two and a half months after our first public release, the Reverie Engine is no longer running through a JavaScript shell. The DSP has been rebuilt as native code on macOS and Windows, and the architecture that drives it has been retuned for long-form ambient work. Under a strict pipeline-parity benchmark, depending on the style, the same 5-minute render runs anywhere from 1.05 to 5.2 times faster, with two to four times less peak memory across the board. The toolbox is also much deeper: the new Custom Chain Builder exposes the full set of 37 DSP modules and lets you arrange them yourself, where 1.0 only offered curated styles.

Product	Reverie
Version	1.13 (released 2026-06-08)
Category	Ambient soundscape generator, standalone desktop app
Platforms	macOS 13+ (Ventura), Windows 10+
Price	EUR 19 launch price (regular EUR 29), one-time payment, lifetime 1.x updates ·
Free tier available	
Modules	37 native DSP modules
Presets	20 Pro presets + 5 Lite presets across 22 style families
Developer	Parallel Minds, Lyon, France
Website	https://reverie.parallel-minds.studio
Contact	William Garrido, Founder · info@parallel-minds.studio

Headline changes

DREAMTIME

From Paulstretch to Dreamtime. Our extreme time-stretch module gets a name of its own. Dreamtime is our evolution of Paul Nasca's 2006 Paulstretch algorithm with six engineering signatures: spectral drift, stochastic multi-segment extraction, phase continuity blending, adaptive segmentation, engine-level duration correction, and seeded phase RNG. The result: live evolving soundscapes, not frozen drones.

NATIVE

Reverie Engine, now native. The DSP pipeline was rewritten end to end in native code on macOS and Windows. Up to 5x faster wall-clock on a 5-min render, 2 to 4x less peak memory across all styles, and no IPC bridge between the UI and the engine.

PRO MODE

Custom Chain Builder. Assemble your own effect chain from the full set of 37 DSP modules. Reorder, tweak parameters, save chains as personal presets. v1.0 only offered curated styles, with no way to compose your own pipeline. Full undo and redo history.

+5 PRESETS

Five new Pro presets. Deteriorating Memories, Harmonic Healing, Forest Breathing, DNA Splice and Ancient Chant join the factory roster. Two new styles, Chimera (spectral DNA fusion) and Sacred (ritual drones), back the last two.

Inside the engine

Architecture choices that survive the rewrite, baked into both native builds.

- **Reproducibility – Composite seed v3.** Chaos, brightness, stretch intensity and the RNG seed are packed into a single 56-bit integer. One number is enough to restore the exact slider state and reproduce a render bit for bit.
- **Determinism – Per-module RNG isolation.** Each module receives its own deterministic sub-seed derived from `hash(main_seed, module_name, chain_position)`. Tweaking one effect never silently shifts the random parameters of another.
- **Headroom – Inter-module soft limiting.** A tanh curve sits between every module to prevent resonance buildup on long chains. Transparent below threshold, audible only when the chain would otherwise clip.
- **Loudness – Long-form LUFS normalization.** Loudness is measured on a 60-second window in the middle of the file (-16 LUFS, +3 dBFS headroom). Adapted to homogeneous drones, lighter on memory than full-file measurement.
- **Architecture – Three-layer preset system.** Module defaults, then style overrides, then factory preset overrides. Recipes never have to mutate a module to fix a single preset, which keeps every other style intact.

DSP highlights, module by module

The nine modules where the rewrite brought a measurable change in either sound quality, numerical stability, or speed.

- **`dreamtime` Six engineering signatures beyond Paulstretch.** Reverie's extreme time-stretch module, now shipped as Dreamtime, is a custom evolution of Paul Nasca's 2006 Paulstretch. Six differences from the baseline: (1) Spectral Drift, read head wanders through the source following three sinusoids at golden-ratio frequencies, drift plus-minus 30%, where the original reads linearly, keeping hour-long renders alive. (2) Stochastic Multi-Segment, for sources over 3 seconds, 2 to 4 non-contiguous micro-segments are extracted and crossfaded, with silent-zone detection and retry, instead of a single contiguous segment. (3) Phase Continuity Blending, 70% new random phase plus 30% previous with shortest-arc wrapping, reducing the grainy digital artefacts of per-frame random while keeping the drone character. (4) Adaptive Segmentation, segment duration auto-scales from 40-70% of the source for short inputs down to 20-40% for longer ones, always guaranteeing at least 40 FFT frames. (5) Duration Correction Loop, engine-level post-chain orchestration with pre-loop, re-stretch and 3.5s sigmoid fades to hit the target duration exactly, where the original is single-pass. (6) Seeded Phase RNG, MT19937 isolated from segment selection RNG, letting you lock the drone texture while varying upstream modules.
- **`spectral_freeze` Per-bin Perlin drift, decoupled phase seed.** Captures multiple phase snapshots and interpolates them slowly. v1.0 jittered each FFT bin with plain random noise; 1.13 replaces it with deterministic 1D Perlin noise (3-octave fbm), so long-form freezes evolve smoothly instead of

shimmering. The phase seed is also decoupled from the texture seed so you can vary one without disturbing the other.

- **temporal_evolution** **Native IIR filtering, no more handcrafted FIR.** Slow macro envelopes (30 to 300 seconds) sweep brightness, cutoff and amplitude to fight drone fatigue. v1.0 reimplemented the filtering with handcrafted scipy windows; 1.13 uses native Butterworth IIR filters, more stable numerically and significantly cheaper per sample.
- **micro_drift** **Double-precision phase accumulator.** Three sub-audible layers (detuning plus/minus 2 to 4 cents, tremolo 0.5 to 2 percent, stereo shift 0.1 to 0.5 ms) modulated by LFOs locked to irrational ratios (ϕ , $\sqrt{2}$). v1.0 computed each LFO sample as $\text{float}(i) * \text{freq} / \text{sr}$, which loses precision past 16 million samples (about 6 minutes at 48 kHz). 1.13 uses a double-precision phase accumulator, which keeps the modulation clean on hour-long renders.
- **spectral_morph** **Proper COLA windowing.** Morphs two spectral snapshots taken at different positions of the source. v1.0 used a simple Hann window for the STFT, which leaves audible overlap-add artefacts; 1.13 ships a proper sqrt-Hann constant-overlap-add (COLA) synthesis window, giving cleaner transitions on the same algorithm.
- **spectral_crosssynth** **Spectral cross-synthesis with COLA cleanup.** A spectral vocoder that imposes the spectral envelope of one part of the audio onto another. Same idea as 1.0; 1.13 inherits the same proper COLA STFT as spectral_morph, which removes the residual smearing of the Python build.
- **polyrhythm_tremolo** **Twelve polyrhythmic schemes, native loop.** African 3:4:5, Fibonacci 3:5:8, Steve Reich 4:5:6, Golden Ratio 1:1.618:2.618 and more. The algorithm is a faithful port of the 1.0 version, but the per-sample loop runs as native code without the Python interpreter overhead.
- **breathing** **In-place asymmetric envelope, single-pass smoothing.** Asymmetric inhale and exhale amplitude envelope with humanized imperfections. v1.0 smoothed the envelope with a Python sliding-window list (allocates each frame); 1.13 does it in place with a single-pass moving average, with much better cache behavior on long files.
- **pitch_drift** **Analog-style wobble, double-precision time-warp.** Slow multi-layer LFO drift that emulates analog or tape pitch instability. The algorithm is the same as 1.0; the gain in 1.13 comes from doing the cumsum-based time-warp in double precision, which avoids the float drift Python accumulated on long renders.

Benchmark methodology

30-minute ambient render, strict pipeline parity on both versions: identical source file, identical preset and parameters, identical mastering chain end to end. Two scenarios because the two main time-stretch families (Dreamtime and spectral_freeze) stress different parts of the engine and gain differently from the native rewrite.

Test conditions

- Hardware: MacBook Air, Apple M3, 8 cores (4 performance + 4 efficiency), 16 GB unified memory, macOS 26.3.1
- Source: 238 s stereo 48 kHz PCM 16 input (electronica)
- Parameters: chaos 70, brightness -50, stretch 35, target 30 min, seed 42
- Mastering chain identical on both versions: DC removal at 5 Hz, multiband resonance taming, brightness-aware high-shelf, soft tanh limiter at 0.95, LUFS normalization to -14 LUFS with -1 dBFS true-peak ceiling, LR2 stereo bass mono crossover at 200 Hz

- Output: stereo 48 kHz WAV PCM 24-bit, around 494 MB per render
- Measured with `/usr/bin/time -l`, file load included in the timing

Dreamtime chain (Dark style)

Chain: *Dreamtime, lowpass, spectral_blur, reverb, breathing*. Stresses time-stretch FFT batching and the IIR filter path.

METRIC	V1.0 (PYTHON)	V1.13 (NATIVE)	DELTA
Wall-clock time	39.08 s	21.87 s	1.79x faster
Peak resident memory	5.10 GB	1.75 GB	2.91x less
User CPU time	35.01 s	38.40 s	native parallelized*
Cycles retired	146.8 G	143.4 G	similar work, shorter wall time
Output WAV duration	1800.3 s	1800.0 s	identical

Spectral freeze chain (Spectral style)

Chain: *lowpass, spectral_freeze, spectral_blur, shimmer_reverb, stereo_width*. Spectral-domain heavy, dominated by STFT/FFT across the chain.

METRIC	V1.0 (PYTHON)	V1.13 (NATIVE)	DELTA
Wall-clock time	29.56 s	24.56 s	1.20x faster
Peak resident memory	5.17 GB	2.79 GB	1.85x less
User CPU time	29.67 s	45.98 s	native parallelized*
Cycles retired	117.8 G	170.6 G	more work in parallel, finishes sooner
Output WAV duration	1781.1 s	1800.0 s	within v1.0 internal 5% target tolerance

* User CPU time is the cumulative work across all cores. The native build is higher because it runs effects in parallel across multiple cores; what you actually wait for as a user is wall-clock time.

Sweep across 8 styles, 5-minute renders

Same source, same parameters, same mastering chain. Only the style changes. Ordered by wall-clock speedup.

STYLE	V1.0 WALL	V1.13 WALL	SPEEDUP	V1.0 RAM	V1.13 RAM	RAM SAVED
Crystalline	35.86 s	6.90 s	5.20x	4.41 GB	1.36 GB	3.24x
Bright	26.55 s	5.49 s	4.84x	4.24 GB	1.14 GB	3.72x
Cloud	14.14 s	6.26 s	2.26x	4.08 GB	1.14 GB	3.58x
Dark	11.63 s	6.36 s	1.83x	4.61 GB	1.13 GB	4.08x
Warm	10.31 s	6.79 s	1.52x	3.22 GB	1.45 GB	2.22x
Deep Drone	10.42 s	7.25 s	1.44x	3.19 GB	1.45 GB	2.20x
Vintage	13.72 s	10.13 s	1.35x	3.62 GB	1.42 GB	2.55x
Spectral	5.91 s	5.62 s	1.05x	2.63 GB	0.91 GB	2.89x

Single-run numbers. v1.13 was measured against the macOS-native Reverie engine, the exact code that ships inside the macOS app. v1.0 was measured against the legacy Python build, with the LR2 stereo bass mono crossover backported

into its mastering pipeline so that both versions run the same mastering chain end to end.

Changelog

Performance

- Engine compiled with AVX2 and FMA on Windows for 256-bit SIMD float operations
- Native Apple Silicon and Intel binaries on macOS
- SIMD-vectorized FFT replaces the Python scipy path on the heaviest spectral modules
- $O(n \log n)$ FFT convolution on reverb and shimmer paths
- Native Butterworth IIR filters in place of handcrafted scipy filters on temporal_evolution and similar modules
- Single-pass in-place processing on modulation modules (breathing, micro_drift), better cache behavior on long renders
- Double-precision phase accumulators on every LFO, fixing float drift that the Python build accumulated past 6 minutes at 48 kHz

Added

- Custom Chain Builder (Pro mode): full set of 37 DSP modules now user-selectable, replacing the curated-styles-only workflow of 1.0
- Five new factory presets: Deteriorating Memories, Harmonic Healing, Forest Breathing, DNA Splice, Ancient Chant
- Multi-format export: WAV, FLAC, OGG, AIFF in 16 or 24-bit
- Composite seed v3 visible in the UI, with full restore of slider state from a seed
- Live seed preview as you move the sliders
- Undo and redo on the entire session
- Multi-select drag and drop
- Native file drag and drop
- Icon picker for custom presets

Improved

- Pixel-perfect native UI on both macOS and Windows
- Faster startup, no JavaScript runtime to warm up
- Tighter audio scheduling, no IPC latency between UI and engine

Fixed

- A class of bridge-related bugs from the 1.0 stack disappeared with the native rewrite

Compatibility

macOS and Windows. Same licenses as 1.0, no re-purchase needed. Output duration unchanged: 30 minutes on Pro, 3 minutes on Free.

Audio samples

Five before/after pairs (original source recordings and their Reverie outputs) are available on the homepage demo section, playable inline:

<https://reverie.parallel-minds.studio/en#demo>

Resources

- Product screenshots: see the `screenshots/` folder in this press kit (3 files, Retina resolution)
- Press page: <https://reverie.parallel-minds.studio/en/press>
- Full changelog: <https://reverie.parallel-minds.studio/en/changelog>
- Release updates (RSS): reverie.parallel-minds.studio/changelog.rss
- Download (free to try): <https://reverie.parallel-minds.studio/en#download>