Grainulator

iOS Edition (iPadOS)

Contents

1. Grainulator

A granular synthesis instrument and effect for macOS, Windows, and iPadOS.

Load any audio sample, slice it into thousands of tiny grains, and reshape those grains into pads, drones, textures, rhythms, and sound effects. Two independent layers, a polyphonic spectral synth per layer, multi-mode reverb, a resonant body model, and a flexible modulation matrix with 5 lanes and 45 destinations.

On first launch, Grainulator loads a factory preset called **First Patch** with a sample already attached — press Play and you're hearing granular synthesis on the second beat.

Quick start

Get sound out of Grainulator in about five minutes.

Read the quick start →

What is Grainulator?

The idea behind the instrument and what you can make with it.

Read the intro →

Parameter reference

Every control on every tab.

Browse the reference →

Troubleshooting

Common issues and how to fix them. See the **Troubleshooting** page in the Desktop or iPad section of the sidebar.

2. What is Grainulator?

Grainulator is a **granular synthesizer**. You give it an audio sample — a recording of anything, from a single piano note to a field recording of rain — and it slices that sample into hundreds or thousands of short fragments called **grains**. Each grain is between about 1 millisecond and half a second long. By spawning grains rapidly, varying their pitch, position, window shape, and amplitude, and layering the result through a modulation matrix, Grainulator turns any source sound into a new instrument.

2.1 Two layers, one engine

Grainulator has two independent **layers**, Layer 1 and Layer 2. Each layer has its own audio buffer, grain engine, window shape, pitch variation, filter, and envelopes. Use the two layers in parallel for thick layered sounds, in opposition for call-and-response patches, or mix a rhythmic layer against a pad for generative backgrounds.

A global synthesis block (resonator, sub oscillator, VCA/filter envelopes, glide) sits downstream of both layers, followed by the FX chain — master filter, compressor, tape delay, and multi-mode reverb. The Layer System reference explains exactly what is per-layer and what is global.

2.2 What you can make with it

- **Pads and drones** from a single held note or a vocal "ahh".
- **Rhythmic textures** by syncing grain density to your DAW's clock and routing a euclidean pattern to the grain trigger.
- **Frozen reverbs** by holding the grain position and sweeping density.
- **Pitched instruments** by mapping the grain source to MIDI notes with the root-note glide control.
- **Sound design** for film, games, and live performance.
- **Cinematic risers and impacts** by modulating scan rate and window shape over time.

2.3 Versions

Grainulator runs as:

- A **standalone desktop app** for macOS and Windows.
- A **plug-in** for any DAW that supports VST3, AudioUnit, or AAX.
- An **AUv3 plug-in** for iPadOS, runnable in any compatible host (AUM, Loopy Pro, Cubasis, GarageBand, Drambo) or as a standalone iPad app.

The interface and sound engine are identical across all versions. Only the installation path, audio routing, and a few input methods differ.

2.4 Where to start

If you want sound right now, read the Quick Start. If you want to learn every control before you touch one, start with the GRAIN tab reference. If you want to build something specific end-to-end, try Your First Patch — a ten-minute walkthrough that ends with a working ambient pad preset.

3. Quick Start

You'll have sound coming out of Grainulator in about five minutes. Pick your platform.

> **First-launch shortcut**
>
> The very first time you launch Grainulator, the app loads a factory preset called **First Patch** with a sample already attached. Press Play (standalone) or hold a MIDI note (plug-in, Instrument mode) and you'll have sound immediately — no setup required. The walk-throughs below assume a clean session with your own sample, but you can also just start tweaking First Patch.

Desktop (Mac / Windows) iPad

1. **Install Grainulator** by following the install guide (see **Desktop → Install** in the sidebar).
2. **Launch the standalone app** (or open Grainulator as a plug-in in your DAW — load it on an instrument or audio track).
3. **Drag any audio file** onto the waveform display at the top of the window. WAV, AIFF, FLAC, MP3, and Ogg are all supported, at any sample rate — Grainulator will resample to the host rate automatically.
4. **Press Play** in the header (standalone), or play a MIDI note from your DAW (plug-in, Instrument play mode).
5. **Turn the Density knob** on the GRAIN tab up to around 30. You should hear a stream of grains pulled from your sample.
6. **Turn the Pitch knob** to ±12 semitones. The grains shift up or down an octave while keeping the same grain size and density.
7. **Open the FX tab** and raise the Reverb Mix knob. Welcome to ambient music.

1. **Install Grainulator** from the App Store (see **iPad → Install** in the sidebar).
2. **Open Grainulator** as an AUv3 plug-in in any compatible host (AUM, Loopy Pro, Cubasis, GarageBand, Drambo) or launch the standalone app.
3. **Tap the waveform display** at the top. The Files browser opens — pick any audio file from iCloud Drive, On My iPad, or your configured cloud storage.
4. **Play a MIDI note** from your host's keyboard, an external MIDI controller, or the standalone app's on-screen keyboard.
5. **Drag the Density knob** on the GRAIN tab upward to around 30.
6. **Drag the Pitch knob** up or down. The grains shift in pitch.
7. **Switch to the FX tab** and raise the Reverb Mix knob.

3.1 Next steps

- For a full guided walkthrough: Your First Patch.
- To understand what you just heard conceptually: What is Grainulator?.
- To learn every control on every panel: start with the GRAIN tab reference and work through the other tabs under **Reference → Tabs** in the sidebar.

4. Your First Patch — An Ambient Pad

This tutorial walks you through building a slowly-evolving ambient pad from a single audio sample. By the end you'll understand the core controls on the GRAIN, TEXTURE, MOD, and FX tabs, and you'll have a sound you can save as a user preset.

Time: about 10 minutes.

You'll need: any sustained sound recording — a held piano note, a vocal "ahh", a string drone, or a field recording with a pitched character all work well. A sharp drum hit is a poor choice for this particular recipe; short percussive samples produce rhythmic textures instead of pads.

4.1 Step 1 — Load a sample

Desktop iPad

Drag your audio file onto the waveform display at the top of the Grainulator window. The waveform appears and Grainulator automatically resamples the file to your current session sample rate on a background thread.

Tap the waveform display at the top of the plug-in. The Files browser opens. Navigate to your audio file (iCloud Drive, On My iPad, or any configured cloud provider) and tap it to load.

You should see a stereo waveform drawn across the top of the window and a vertical line marking the current grain spawn position.

4.2 Step 2 — Set the grain core

Open the **GRAIN** tab if it isn't already visible and dial in these starting values in the GRAIN column:

- `Size` — set to around **0.3 s** (300 ms). Large grains are the foundation of pad sounds: they play back enough of the source that each grain has a recognizable tone, but not so much that you just hear the original sample looping.
- `Density` — set to around **20**. Enough grains overlap to create a continuous texture without the CPU load of hundreds per second.
- `Pitch` — leave at **0** for now. We'll add variation with the MOD tab later.
- `Speed` — set to around **0.05**. This is the *playback* speed of the grain reader through the buffer — slow enough that successive grains come from nearly the same point in the sample, which is what makes a pad sound static.
- `Jitter` — set to about **0.15**. A small amount of random position offset prevents the machine-regularity that pure 20 Hz density would otherwise produce.

Press Play (standalone) or hold a MIDI note (plug-in, Instrument mode) and you should already hear a rough pad-like sound.

4.3 Step 3 — Pick a window shape

Switch to the **TEXTURE** tab. In the **Shape** column:

- Set `Window Shape` to **Hanning**.

The window is the amplitude envelope applied to every grain. Hanning (a raised-cosine curve) fades each grain in and out smoothly, which is exactly what you want for pad sounds — the fade-in and fade-out eliminate the clicks you would get from abrupt grain starts, and the soft envelope lets adjacent grains blend rather than compete. Try switching to **Triangle** or **Trapezoid** for comparison: they sound noticeably more percussive.

Leave Tilt, Curve, and Sides at their defaults for now.

4.4 Step 4 — Add motion with modulation

Pads get boring when nothing moves. Switch to the **MOD** tab and set up one of the five mod lanes:

- `Source` — **LFO**.
- `Dest` — **Position** (grain spawn position).
- `Amount` — about **0.15**. This is a bipolar depth, so the grain position will drift ±15% of the buffer around the currently-set point.
- `LFO Rate` — about **0.1 Hz** (one full cycle every 10 seconds).
- `LFO Shape` — **Sine**.

The position of the grains now slowly drifts through the sample over time. Because your sample is a sustained sound rather than a percussive one, the drift sweeps through slightly different timbres (different overtones, slightly different decay characteristics) and adds the organic evolution that a static pad lacks.

4.5 Step 5 — Send it through reverb

Switch to the **FX** tab. In the Reverb section:

- `Mode` — **Stormy**. Of the three reverb modes, Stormy has the longest tail and the most modulated character, which is the ambient-pad default.
- `Size` — **0.75**. A large virtual space.
- `Damping` — **0.4**. Some high-frequency absorption so the reverb tail doesn't get harsh.
- `Mix` — **0.45**. A generous wet amount — ambient pads live and die by their reverb.

If you want to push into washing-drone territory, raise Mix to 0.7 and enable `Freeze` on the reverb to hold the current tail indefinitely.

4.6 Step 6 — Save it as a preset

In the header bar, click the **Save** button (disk icon next to the preset name). Enter a name — "My First Pad" is traditional — and confirm. The preset file is written to your user presets directory and is now available from the preset browser on any future launch.

4.7 What you just learned

- The four most important grain parameters: **Size**, **Density**, **Speed**, and **Position**. Every pad recipe starts with these.
- How **window shape** changes the *character* of grains beyond their basic parameters — a Hanning window is smooth, a Trapezoid window is punchy, and the choice matters as much as any knob on GRAIN.
- How **modulation** adds movement to an otherwise static patch, and how even a very slow LFO on a single parameter can make a sound feel alive.
- How **reverb mode** is more than a tail length — it's a tonal choice, and the Stormy mode sounds different from Misty even at identical Size and Mix values.

4.8 Where to go next

- Read the full GRAIN tab reference to discover every control you didn't touch (including the MORPH column's assignable Macro knob, which is a tutorial of its own).

- Read the full MOD tab reference — the single lane you set up is one of five, and the rhythmic sources can drive entirely different types of patches.
- Try the same recipe with a **percussive** sample (a single drum hit, a snippet of speech, a guitar pluck). The same knob positions will produce a completely different sound, which is the whole point of granular synthesis.

5. What is granular synthesis?

5.1 Opening hook

Imagine you have a vinyl record playing a single piano note. The note holds steady for thirty seconds. Now imagine you pick up the needle, drop it back on the same groove, wait a few milliseconds, and repeat — ten times a second. You'll hear a rapid stutter. The original piano note is in there, but the sound itself is something new: a buzzing, textured drone stitched together from tiny snippets of the original.

That's granular synthesis. The technique takes any recorded sound, slices it into fragments shorter than a human ear can resolve individually, and resynthesizes a new sound out of the overlap.

5.2 The core idea

A **grain** is a short audio fragment — typically between one millisecond and half a second long — taken from a source recording and shaped by a smooth amplitude envelope at both ends so it doesn't click when it starts or stops. By itself a single grain sounds like what it is: a tiny snippet of sound. The interesting part starts when you spawn many grains in rapid succession, each from a slightly different position in the source buffer, and let them overlap. Your ear can't resolve individual grains below about twenty per second, so above that rate the grains fuse into a continuous texture.

Four variables define how that texture sounds. **Grain size** — how long each grain is — controls whether you hear something pad-like (longer grains, more of the original sample audible in each) or percussive (shorter grains, each heard as a pointillistic click or pluck). **Density** — how many grains spawn per second — controls whether the result sounds sparse and rhythmic (low density, grains heard as discrete events) or smooth and continuous (high density, grains overlapping into a wash). **Position** — where in the source audio each grain is sliced from — lets you hold a single moment of the source or scan through it at any speed you like, including backwards. **Pitch** — the per-grain transposition — lets you shift the perceived pitch of the resynthesized sound without changing how fast the position scans, because each grain is short enough that the pitch-shift is effectively instantaneous.

Those four knobs are the heart of every granular patch ever made. Every other granular control — jitter, spread, window shape, trigger patterns, modulation — is a variation or a decoration on top of them.

The technique was first theorized by the physicist Dennis Gabor in 1947, in a paper called *Theory of Communication*. Gabor wasn't a composer; he was working on a mathematical decomposition of signals into elementary time-frequency atoms. Musically, the technique was first put into practice by Iannis Xenakis in the late 1950s — he assembled grain sequences by hand, cutting and splicing magnetic tape, because no computer of the era could do it in real time. Curtis Roads wrote the software implementations that made granular synthesis practical on computers in the 1970s and 1980s, and his book *Microsound* (2001) remains the canonical reference.

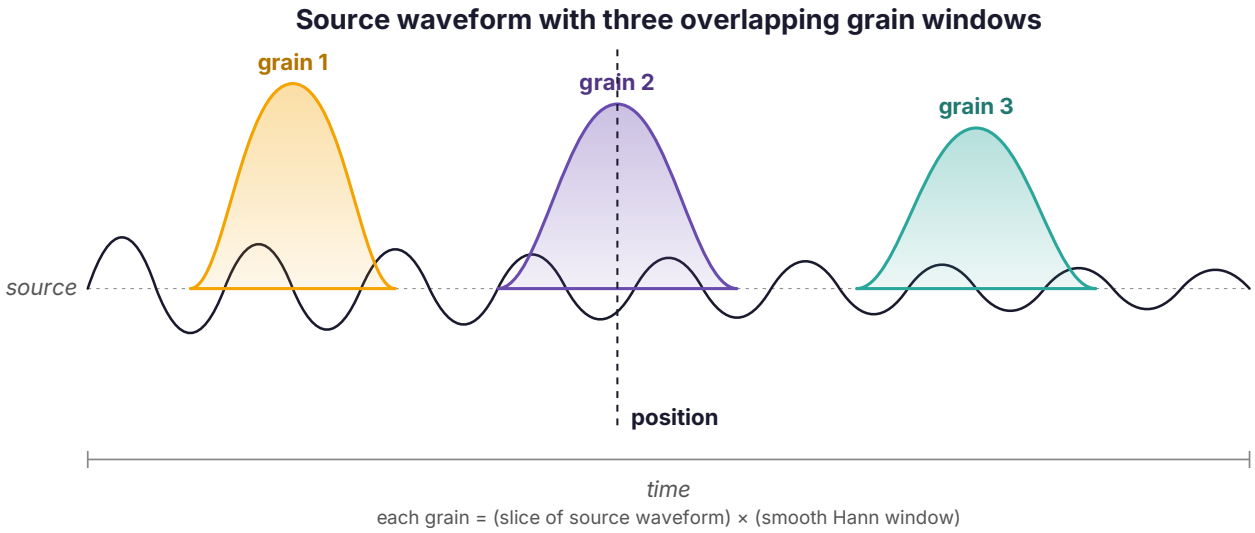

The diagram above shows a source waveform (the piano note) with three grain windows overlaid at different positions. Each grain takes a small slice of the source and multiplies it by a smooth envelope that fades in and out. The grains overlap each other, and the listener's ear blends them into a continuous texture.

5.3 In Grainulator

Everything above maps directly to controls on the GRAIN tab:

- `Size` sets the grain length, from 1 ms (so short you hear each grain as a click or as its own pitch) to 500 ms (so long it sounds almost like the original sample playing). Most musical pad textures live in the 30–200 ms range.
- `Density` sets how many grains spawn per second. At low values (below about 10 Hz) you hear discrete grain events — useful for rhythmic, pointillistic patches. Above about 20 Hz the grains fuse into a continuous wash; above about 100 Hz the grain-spawn rate itself starts to be heard as pitch.
- `Pitch` transposes each grain up or down by a number of semitones at the moment it spawns. Because grains are short, you can pitch-shift them dramatically without the time-stretching artifacts you'd hear from pitch-shifting the entire sample in real time.
- The position cursor on the waveform display sets where in the source audio each grain is sliced from. Hold the cursor still and you get a "frozen" sound at one moment of the source. Sweep it with the `Speed` knob, or modulate it from the MOD tab, and you scan through the source in time.

The TEXTURE tab adds the window shape that each grain is multiplied by — Hanning for smooth pads, Trapezoid or Pluck for percussive attacks — plus pitch-variation and scan controls that modulate the core four parameters over time. The MOD tab lets you route any of those parameters to an LFO, envelope, or clock-synced rhythmic source to put the whole patch into motion.

The math, for the curious

A grain is a finite-length signal $g(n)$ formed by multiplying a source signal $s(n)$ by an envelope window $w(n)$:

$$g(n) = s(n_0 + n) \cdot w(n), \quad 0 \le n < L$$

where L is the grain length in samples and n_0 is the read position in the source buffer. The window function $w(n)$ is typically a Hann or Tukey window — smooth at both ends so that the product of source-times-window fades in and out without an audible click. The output of the granular synthesizer is the sum of many overlapping grains, each spawned at its own start time t_i:

$$y(n) = \sum_i g_i(n - t_i)$$

Density is the rate λ at which spawn times t_i are scheduled — either uniformly (every $1/\lambda$ seconds) or stochastically (Poisson-distributed around a mean rate of λ). **Pitch shifting** is implemented by reading the source buffer at a different speed inside each grain: a per-grain resampling that doesn't affect the overall playback rate, because each grain is short enough that the resampling is inaudible as a time stretch.

Curtis Roads' *Microsound* (MIT Press, 2001) derives all of this formally, connects it to the Fourier-style time-frequency decomposition Gabor originally proposed, and is the canonical reference for anyone who wants to dig further.

5.4 Further reading

- **Microsound**, Curtis Roads (MIT Press, 2001) — the definitive book on granular synthesis theory and practice.
- **Theory of Communication**, Dennis Gabor (Journal of the Institution of Electrical Engineers, 1946) — the original mathematical paper that introduced the grain concept.
- **Formalized Music**, Iannis Xenakis (Indiana University Press, 1971) — the first composer to put the theory to musical use writes about it in his own words.

5.5 See also

- GRAIN tab reference
- Your First Patch — apply this theory in ten minutes.

6. Header Bar

The header bar runs across the top of the Grainulator window. It holds the preset browser, transport, host-sync toggle, and (in compact layouts) the layer selector and tab navigation.

6.1 Preset browser

At the center of the header bar is the **preset browser**:

- **← / →** — step to the previous or next preset.
- **Name** — click to open a searchable list of factory and user presets.
- **Save** — write the current patch as a user preset (`.gv2p` file).
- **Init** — load the default "init" preset.

User presets are stored in the application support directory and are shared between the standalone app and the plugin formats.

6.2 Transport

- **Play / Stop** — toggles global playback. Behavior depends on the play mode (see below).
- **SYNC** — a toggle that locks grain density, LFO rate, and delay time to the host's tempo. In the standalone app the toggle is disabled; in a DAW it follows the host transport.

6.3 Play mode

Grainulator has two play modes:

- **Free** — grains spawn continuously as long as a layer is enabled. The Play button drives playback.
- **Instrument** — grains spawn only while MIDI notes are held, and each note sets the root pitch for the layer.

The play-mode toggle is available from the header settings menu and from any MIDI-control panel on the SYNTH tab.

6.4 Layer selector

In compact layouts (AU plugin and iPhone), only one layer's controls are visible at a time. A **Layer 1 / Layer 2** segmented control in the header (or at the top of the GRAIN panel in the full layout) switches the visible layer. The selector is highlighted in the current layer's accent color.

Internally, Grainulator calls these "voices" (parameter names start with `v1` or `v2`), but the UI uses "Layer" everywhere a user sees them.

6.5 Tab navigation

The **GRAIN**, **TEXTURE**, **SYNTH**, **FX**, and **MOD** tab buttons sit beneath the header bar. Click a tab to switch the main panel view. The MOD tab is always visible, even in layouts that collapse other panels, so the modulation matrix is reachable from anywhere.

6.6 See also

- Waveform Display
- Layer System

7. Waveform Display

The waveform display sits beneath the header bar and shows the loaded audio sample together with the current grain activity. In the full (two-row) layout there are **two** waveform displays side by side — one for each layer. In compact layouts the selected layer's display is shown full width.

7.1 What you see

- **Stereo waveform** — the loaded sample drawn as a pair of channels.
- **Position cursor** — a vertical line marking the current grain spawn position in the buffer.
- **Active grain dots** — each currently sounding grain is drawn as a small dot over the waveform. The horizontal position encodes the grain's position in the buffer; the vertical position encodes pan; the brightness encodes amplitude.
- **Trim handles** — a pair of bars at the left and right edges of the audible region. Hover the waveform to reveal them; drag to set the loop in / out points. Double-click a handle to reset that side to the buffer edge. The handles auto-reset when you load a new sample or enter Live recording mode.
- **Freeze button** — captures the current buffer state so you can scroll the position freely without the playhead advancing (useful for holding a specific moment for manipulation).
- **Layer-link toggle** — on Layer 2's waveform display only, a chain icon toggles whether Layer 2 uses its own buffer (independent) or shares Layer 1's buffer.

7.2 Loading a sound

You can load an audio file three ways:

1. **Drag and drop** an audio file onto the waveform display.
2. **Click** anywhere on the waveform display to open a file browser.
3. **Use the Load menu** in the header (in full layout) or the layer's context menu (right-click / long-press).

Supported formats include WAV, AIFF, FLAC, MP3, and Ogg Vorbis. Files at any sample rate are automatically resampled to the current app or host sample rate on a background thread, so loading a 48 kHz file into a 44.1 kHz session plays back at the correct pitch.

Clear Audio

The per-layer context menu (right-click on a waveform display, or long-press on iPad) includes a **Clear Audio** item. Selecting it empties that layer's buffer and resets its trim handles. Useful when you want to disarm a layer without unloading the preset.

7.3 Setting the position

Click and drag on the waveform to scrub the grain spawn position manually. In Free play mode the position advances according to the Speed parameter; drag sets it explicitly. In Instrument play mode, the position advances only while notes are held.

7.4 See also

- Header Bar
- Layer System

8. Layer System

Grainulator has two independent **layers**, Layer 1 and Layer 2. Each layer has its own audio buffer, grain engine, window shape, pitch variation, filter, resonator, sub oscillator, and envelopes. The two layers share the global effects chain (delay, reverb, compressor, master filter) and the modulation matrix.

Internally the code calls these "voices" and parameter names begin with `v1` / `v2`, but the UI and this manual use "Layer 1" / "Layer 2" throughout.

8.1 What each layer owns

Truly per-layer controls — independent between Layer 1 and Layer 2 — live on the GRAIN and TEXTURE tabs:

- Grain core (size, density, pitch, speed, jitter, spread, pan, level)
- Window shape, window morph, decay
- Pitch variation (probability, mode, direction, range)
- Dynamics randomisation
- Reverse probability
- Scan rate / depth
- Loop points, freeze state
- The per-layer filter (model, cutoff, resonance, drive, env depth)

On the SYNTH tab, only `Root 1` and `Root 2` are truly independent. The Resonator, Sub Oscillator, VCA envelope, Filter envelope, and Glide are all **global**: they are mounted in the UI via Layer 1's parameters, and any change applies to both layers. This is a v2-era simplification — the per-layer parameter slots still exist in the preset format for backward compatibility, but the audio engine runs a single shared instance.

See the SYNTH tab page for details.

8.2 What is shared between layers

The following are global and affect both layers together:

- SYNTH panel: resonator, sub oscillator, VCA / filter envelopes, glide
- FX chain: master filter, compressor, delay, reverb
- Master gain, dry/wet
- Modulation matrix (each lane can be targeted at a specific layer)
- Macro knob (the hero Morph control on the GRAIN tab)

8.3 Buffer modes

A **Link** toggle on Layer 2's waveform display controls whether the two layers share one audio buffer or use independent ones:

- **Independent** (unlinked) — each layer loads its own sample. Dropping a file onto Layer 1's waveform fills Layer 1; dropping onto Layer 2's waveform fills Layer 2.
- **Linked** (shared) — Layer 2 reads from Layer 1's buffer. Loading a new sample into Layer 1 updates both layers. Useful for building dense textures from a single source.

In linked mode, Layer 2 still has its own grain parameters — it's only the underlying audio that is shared.

8.4 Enabling and muting layers

Each layer has a power toggle that disables it entirely (no grains spawn, no CPU spent). Mute and solo buttons provide quick in-session silencing without resetting the layer.

8.5 Switching layers in compact mode

In the full two-row standalone layout, both layers' controls are visible side by side. In the compact AU or iPhone layout, only one layer's controls are visible at a time and the **Layer 1 / Layer 2** segmented control in the header bar switches the view. Changing layers never stops playback — both layers keep running, you are only changing what is visible for editing.

8.6 See also

- Header Bar
- Waveform Display

9. GRAIN Tab

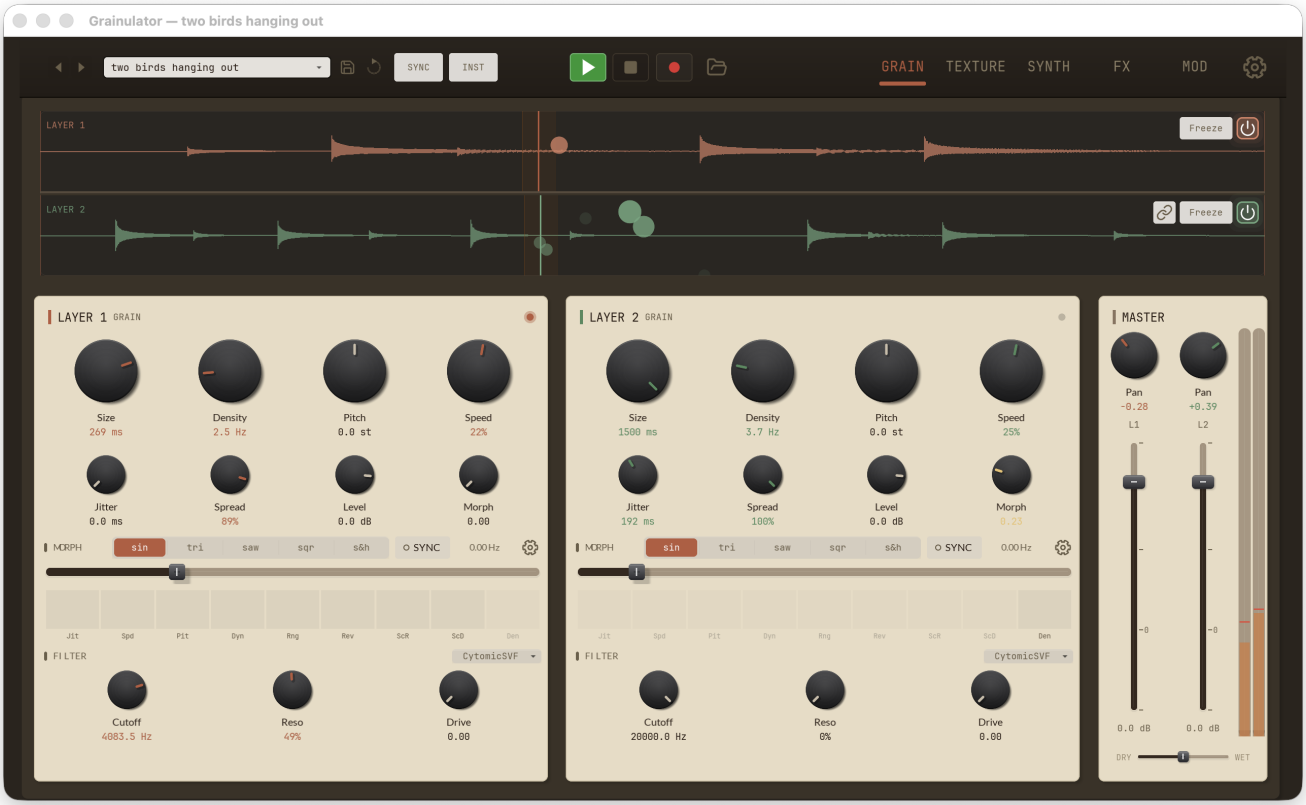

The GRAIN tab is Grainulator's default view and the heart of the granular engine. The desktop layout shows **both layers side by side, all the time** — there is no layer-focus toggle. From left to right the tab is:

- **LAYER 1** panel — Layer 1's grain core, per-layer Morph, and per-layer filter, stacked top-to-bottom.
- **LAYER 2** panel — the same controls for Layer 2, independently set.
- **MASTER** panel — pans, faders, the global VU meter, and the Dry/Wet blend.

A pair of waveform displays sits above the panel row, one per layer. See Waveform Display for everything they do.

9.1 Per-layer panels

Layers 1 and 2 carry the same controls; setting them independently is how you build dual-source patches, call-and-response textures, and layered pads. Each panel has three sub-sections — **GRAIN**, **MORPH**, and **FILTER** — stacked from top to bottom.

GRAIN sub-section

Eight knobs in two rows of four drive the granular engine.

Control	Range	Description
`Size`	0.001–1.5 s	Duration of each grain. Short sizes give choppy, rhythmic textures; longer sizes produce smooth, pad-like tones. Latched at grain spawn.
`Density`	1–512 Hz	Grain spawn rate. Higher values create denser, overlapping clouds; lower values leave sparse, audible gaps.
`Pitch`	−24 to +24 st	Transposes each grain in semitones. Applied at spawn time, so moving it re-pitches only new grains.
`Speed`	−3 to +3 (displayed as %)	Internal playback speed multiplier. Negative values play in reverse; values outside ±1 combine with `Pitch` for extreme time-stretch effects. Displayed as a percentage (100 % = 1.0). Latched at grain spawn.
`Jitter`	0–0.5 s	Random offset added to each grain's start position. Small amounts humanize the stream; larger amounts scatter grains across the buffer.
`Spread`	0–1	Stereo spread of grain pan. 0 is mono, 1 is full L/R separation.
`Level`	−70 to +12 dB	Per-layer grain level. Mirrors the layer fader in the MASTER panel — moving one moves the other.
`Morph`	0–1	Per-layer macro that simultaneously moves every assigned morph target. 0 leaves parameters at their base values; 1 drives them to their configured target values. Targets and depths are configured from the gear button in the MORPH strip.

`Size` and `Density` interact: their product is roughly the overlap. At 100 Hz with 50 ms grains you have ~5 grains at once — a continuous drone. Drop density to 5 Hz and you hear each grain individually. `Jitter` then breaks up machine-gun periodicity.

MORPH strip

Below the eight grain knobs sits the per-layer **Morph LFO** strip — an LFO that automates the `Morph` knob above. The two together are how a single layer evolves over time without any MOD lane assignment.

Control	Range	Description
`ON`	toggle	Enables the Morph LFO for this layer. Off = the `Morph` knob is static.
`Wave`	Sine / Triangle / Saw / Square / S&H	LFO shape, picked from the inline wave-selector buttons.
`SYNC`	toggle	When on, the rate snaps to the host tempo at the configured division. When off, `Rate` is in Hz.
`Rate`	0–20 Hz (free) / 1/32 – 1/1 (synced)	LFO speed. Very slow rates (below 0.1 Hz) work well for long evolving pads.

A horizontal bar beneath the controls visualises the live LFO output across each assigned target.

The gear button at the top right of the MORPH strip opens the morph settings popup, where you assign which parameters the `Morph` knob drives and at what depth.

FILTER sub-section

The FILTER row is the **per-layer** filter, upstream of the global master filter on the FX tab. Only the tone-shaping controls live here; the filter envelope (ADSR and Env Depth) is on the SYNTH tab.

Control	Range	Description
`Filter Model`	10 models	Dropdown selector: Stilson, Microtracker, Krajeski, MusicDSP, Oberheim, Improved, RKSimulation, Hyperion, DaisyLadder, CytomicSVF. Each model has a distinct saturation and resonance character. Defaults to CytomicSVF.
`Cutoff`	20 Hz – 20 kHz	Filter cutoff frequency.
`Reso`	0–1	Resonance amount at the cutoff. High settings self-oscillate on some models.
`Drive`	0–1	Input drive into the filter stage. Adds saturation and harmonic content, especially at higher resonance.

Layer header

Each layer panel's header shows **LAYER 1 GRAIN** / **LAYER 2 GRAIN** and a small activity indicator. The indicator lights when any grain is sounding on that layer — a quick way to verify a layer is producing audio without watching the meters.

9.2 MASTER panel

The Master panel sits to the right of the two layer panels. It moved here from the FX tab in the 2026-04-07 layout revision so layer balance and wet/dry are always visible on the default view.

Control	Range	Description
`Layer 1 Pan`	−1 to +1	Stereo pan for Layer 1.
`Layer 2 Pan`	−1 to +1	Stereo pan for Layer 2.
`Layer 1 Level`	−70 to +12 dB	Layer 1 vertical fader. Mirrors the `Level` knob in Layer 1's GRAIN row.
`Layer 2 Level`	−70 to +12 dB	Layer 2 vertical fader. Mirrors the `Level` knob in Layer 2's GRAIN row.
Master VU meter	—	Stereo meter on the master output (post layers, post wet/dry).
`Wet/Dry`	0–1	Global wet/dry balance between the granular output and the dry input.

The faders carry small tick marks at common reference levels (−∞, −12, −6, 0, +6) so you can balance two layers by feel.

9.3 Dual waveform display

Above the panel row, the GRAIN tab draws **two waveform displays side by side**, one per layer, each with its own position cursor, active-grain dots, freeze button, and trim handles. Layer 2's display also carries the layer-link toggle for switching between shared and independent buffers. See Waveform Display for the full treatment.

9.4 See also

- Waveform Display
- Layer System
- TEXTURE Tab
- SYNTH Tab — Spectral, Voice Interaction, resonator, sub, envelopes
- FX Tab

10. TEXTURE Tab

The TEXTURE tab shapes the *character* of each grain once the GRAIN tab has decided how big, how dense, and from where in the buffer it reads. Everything here is per-layer — Layer 1 and Layer 2 each have their own independent TEXTURE settings. The tab is organised into five columns: **Shape**, **Pitch**, **Dynamics**, **Reverse**, and **Scan**.

10.1 Shape — Grain Window

Every grain has an amplitude envelope ("window") applied on playback. The window determines whether grains sound smooth and pad-like, punchy and percussive, or plucked and decaying. A live preview above the knobs shows the current shape; tweaking Tilt, Curve, or Sides updates it in real time.

Control	Range	Description
`Window Shape`	Hanning / Gaussian / Trapezoid / Triangle / Tukey / Pluck / PluckSoft / ExpDecay	Base envelope shape applied to every grain.
`Tilt`	0 – 1	Shifts the envelope peak earlier or later within the grain. 0.5 is symmetric.
`Curve`	0 – 1	Bends the attack and release slopes between linear and exponential.
`Sides`	0 – 1	Softens or sharpens the edges of the window (trapezoid / Tukey plateau width).
`Decay`	1 – 10	Exponent for the Pluck / PluckSoft / ExpDecay shapes. Higher = faster fall-off.

Hanning and Gaussian give the smoothest, most classical granular sound. Trapezoid and Triangle produce punchier, more rhythmic grains. Pluck, PluckSoft, and ExpDecay give each grain a struck-string envelope that is great for pitched percussion and pointillist textures.

Morph

A **MORPH** toggle lets the window crossfade between two shapes over time. When enabled, a second dropdown (Window B) and two extra knobs appear.

Control	Range	Description
`Window B`	same 8 shapes as Window Shape	Target window for the morph.
`Morph`	0 – 1	Static crossfade position between Window A and Window B.
`Rate`	0 – 20 Hz	Automatic morph LFO speed. 0 = static (use the Morph knob manually), higher values cycle between A and B.

10.2 Pitch Variation

Pitch Variation re-tunes individual grains at spawn time to add harmonic richness, arpeggios, or scale-snapped generative melodies. It stacks on top of the base Pitch knob on the GRAIN tab. Only grains that "win" the Probability dice roll are re-pitched.

Control	Range	Description
`Prob`	0 – 100 %	Chance that a new grain is re-pitched. 0 % = always use base pitch, 100 % = every grain.
`Mode`	Octaves / Oct+5ths / Scale	Octaves picks random octave transpositions; Oct+5ths adds perfect fifths; Scale snaps grains to a musical scale.
`Direction`	Both / Up / Down	Constrains variation to move only upward, only downward, or in either direction from the base pitch.
`Range`	1 Oct / 2 Oct / 3 Oct	Maximum transposition distance.

In Scale mode an additional scale preset (0–11) selects the quantisation scale used to snap re-pitched grains. The preset is set in the pitch-variation engine rather than from a top-level knob.

Note: **Glide** (pitch smoothing between notes) was previously shown here but is now on the SYNTH tab's MIDI panel.

10.3 Dynamics

Dynamics randomises each grain's amplitude at spawn, adding natural breathiness or aggressive volume jitter depending on the Range setting. The base grain level itself lives on the GRAIN tab (`Level`); this section only adds variation on top of it.

Control	Range	Description
`Probability`	0 – 100 %	Chance that a grain's amplitude is randomised.
`Range`	0 – 1	Maximum amplitude deviation. 0 = no change, 1 = grains can be fully attenuated.

10.4 Reverse

Control	Range	Description
`Probability`	0 – 100 %	Chance that a grain plays backward. 50 % gives an even mix of forward and reversed grains; 100 % reverses every grain.

Because only the *grain's* playback direction flips — not the source buffer — Reverse works with any position, speed, or loop setting and can be modulated for rhythmic forward/backward effects.

10.5 Scan

Scan applies a slow position LFO to the grain spawn point, sweeping through the buffer over time. Combine with a static Position knob on the GRAIN tab to drift around a focal point, or with Speed 0 to replace static freezing with gentle motion.

Control	Range	Description
`Rate`	0 – 5 Hz	Speed of the scan LFO. 0 disables scanning.
`Depth`	0 – 0.5	Amount of buffer traversed by the LFO, as a fraction of the whole file. 0.5 covers half the buffer peak-to-peak.

10.6 See also

- GRAIN Tab
- SYNTH Tab
- Layer System

11. SYNTH Tab

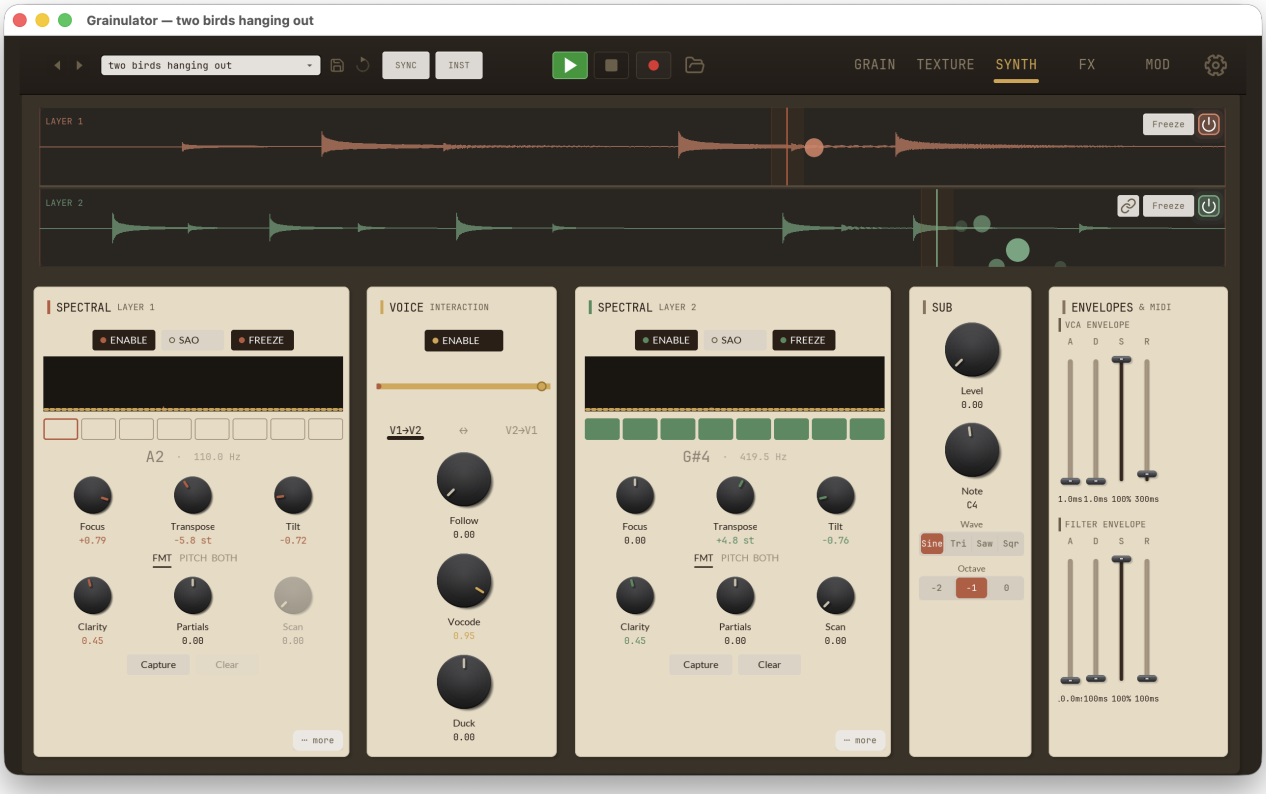

The SYNTH tab gathers the traditional-synthesis elements that sit on top of the granular core: the two **Spectral Array Oscillator (SAO)** columns — per-layer phase-vocoder synths — a **Voice Interaction** panel between them for cross-layer routing, a **Sub Oscillator**, and the **VCA and Filter envelopes** that shape each played note. The **Resonator** lives on the SYNTH tab in compact layouts and elsewhere in the desktop layout; see the section below.

Unlike the GRAIN tab — where every knob edits one layer at a time — most of the SYNTH panels drive a single shared engine. The Sub Oscillator, VCA envelope, Filter envelope, and Glide are **global**: they are mounted in the UI via Layer 1's parameters and any change applies to both layers. The two **Spectral** columns are genuinely per-layer — Layer 1 and Layer 2 each have an independent Spectral Array Oscillator. Only `Root 1` and `Root 2` under MIDI are genuinely independent per layer.

11.1 Spectral panel (Spectral Array Oscillator)

Each layer has its own **Spectral Array Oscillator (SAO)** column that sits at the head of the SYNTH tab — Layer 1's panel on the left, Layer 2's on the right, with the Voice Interaction panel between them. The SAO is a phase-vocoder-based synth that captures or generates a spectral frame and then re-tunes, re-shapes, and re-pitches it. It is **MIDI-responsive and polyphonic** — held notes drive multiple voices through a single SAO instance.

Top toggles

A row of three toggles runs along the panel header:

Toggle	Behaviour
`ENABLE`	Master power for the Spectral panel. Off bypasses the SAO entirely and consumes no CPU.
`SAO`	Switches between **Spectral Array Oscillator** mode (synth-style polyphonic re-pitch driven by MIDI) and the simpler frame-pass mode. SAO mode is the default for melodic playing.
`FREEZE`	Captures the current spectral frame and holds it. Useful for sustaining a moment of the source under polyphonic re-pitch.

8-step pitch grid

Beneath the toggles, an **8-step pad row** displays the SAO's pitch sequence — each cell shows where in the spectral frame the next voice will read. The grid scrolls under polyphonic activity; the active cell is highlighted in the layer's accent color (orange for Layer 1, green for Layer 2). The current center pitch (e.g. `A2`, `G#4`) and its frequency are displayed alongside the grid.

Knobs (top row)

Three knobs above the FMT toggle shape the spectral frame's pitch and balance:

Control	Range	Description
`Focus`	−1 to +1	Sharpens or softens which partials are emphasised in the spectral frame. Negative values broaden the spectrum; positive values concentrate energy on the strongest partials.
`Transpose`	−24 to +24 st	Pitch offset applied to the entire spectral frame in semitones. Independent of MIDI note tracking.
`Tilt`	−1 to +1	Tilts the spectrum: negative darkens (rolls off highs), positive brightens (rolls off lows).

FMT toggle

A `FMT — PITCH BOTH` segmented toggle sits between the two knob rows. It controls how MIDI notes interact with the spectral frame:

State	Behaviour
`FMT`	Formant-only — incoming notes shift formants without changing fundamental pitch.
`PITCH`	Pitch-only — notes re-pitch the frame without preserving formants.
`BOTH`	Both formant and pitch track the played note. The default — the most "synth-like" behaviour.

Knobs (bottom row)

Control	Range	Description
`Clarity`	0–1	Removes noise and softens artifacts in the spectral frame. Higher values produce a cleaner, more pitched tone at the cost of texture.
`Partials`	0–1	Number of harmonic partials retained, as a fraction of the available bins. Lower values produce thinner, more focused tones; higher values include more harmonic detail.
`Scan`	0–1	Position scan through the spectral frame buffer. At 0 the SAO reads the captured frame; turning Scan up sweeps through earlier frames in the buffer for textural movement.

Capture / Clear

Two buttons at the bottom of the panel control the spectral frame:

- `Capture` — grabs a snapshot of the live grain output as the new spectral frame. Press while audio is playing through the layer to load fresh material into the SAO.
- `Clear` — wipes the captured frame and returns to silence until the next capture.

The captured frame is per-layer and is saved with the preset.

11.2 Voice Interaction panel

Sits between the two Spectral columns. Replaces the former **MASTER · XMOD** panel in v20+. Three knobs drive categorically different cross-voice effects — **Follow** (time-domain source coupling), **Vocode** (spectral envelope transfer), and **Duck** (amplitude sidechain) — governed by a shared direction switch and a power toggle that bypasses the whole section.

An animated bus indicator above the direction switch visualises current effect amount (line thickness) and routing (flow direction of the travelling dots). In **Both** mode the dots follow whichever layer is louder at the moment.

Direction

One 3-state segmented control governs all three knobs:

State	Behaviour
`V1 → V2`	Layer 1 is driver; Layer 2 is follower. All three knobs affect Layer 2 only.
`V1 ↔ V2` (Both)	Mutual / symmetric. Each layer acts on the other with the same knob amounts.
`V2 → V1`	Layer 2 is driver; Layer 1 is follower.

Knobs

Control	Range	Description
`Follow`	0–1	The follower layer's granular engine reads the driver's **live post-grain audio** instead of its own loaded buffer. At 1 the follower is fully tracking the driver's output; at 0 it plays its own buffer (default). The follower's `Position` knob is reinterpreted as ring-buffer age — 0 = newest, 1 = oldest (~2 s back). Grain size, pitch, and speed still work inside the live source. In Both mode, a one-block latency prevents feedback.
`Vocode`	0–1	True spectral vocoding. The driver's mean-normalised spectral envelope multiplies the follower's bin magnitudes. At 1 the follower adopts the driver's timbral/formant contour while keeping its own harmonic content. Loudness-preserving — the mean normalisation stops the follower from getting quiet when the driver is sparse. This replaces the old X-Mod enhancer, which was a subtle per-bin boost; Vocode is categorically different.
`Duck`	−1 to +1	The driver's amplitude envelope modulates the follower's output gain. Negative = classic sidechain duck (follower dips when driver is loud). Positive = "swell" / shadow (follower amplifies with the driver). 0 is bypass. Attack ≈ 3 ms, release ≈ 150 ms, hard-coded in v20. Read from pre-duck envelopes so mutual ducking in Both mode is feedback-safe.

Power toggle

The power icon in the panel title stripe hard-bypasses the entire section. Use it to A/B the contribution of Voice Interaction without resetting all three knobs.

Mod routing

All three knobs are available as mod destinations under a new **Voice FX** category on the MOD tab. The direction switch and power toggle are not modulatable (they're discrete state), but their param values round-trip cleanly through preset save/load and DAW automation.

Preset upgrade from v19 and earlier

Presets saved in v19 with a non-zero X-Mod value are automatically upgraded on load: **Vocode** is seeded to `max(V1 X-Mod, V2 X-Mod)`, direction is set to **Both**, and the section is enabled. The legacy per-layer X-Mod params are retained in the enum as no-ops so old mod-lane routings do not produce errors — their destination menu entry now reads `(legacy)`.

11.3 Resonator

A global resonant body based on the Mutable Instruments *Rings* DSP. When `Enable` is off the resonator is bypassed and consumes no CPU. A small live preview next to the panel visualises the current harmonic structure.

Control	Range	Description
`Enable`	on / off	Bypass the resonator entirely. Default off.
`Polyphony`	1 / 2 / 3 / 4	Number of simultaneous resonator voices. Higher values give richer chords at higher CPU cost.
`Chord`	0–1	Harmonic structure. Sweeps through chord qualities and inharmonicities (internal name: *Structure*).
`Bright`	0–1	Brightness — the balance of upper partials. Higher values add treble and shimmer.
`Damp`	0–1	Damping / decay time of the resonance. Low = short ping, high = long ring.
`Scatter`	0–1	Modal position / excitation point. Shifts the harmonic balance (internal name: *Position*).
`Note`	0–127 (MIDI)	Center pitch of the resonator. Displayed as a note name.
`Mix`	0–1	Dry/wet blend between the pre-resonator signal and the resonator output.
`Distort`	0–1	Drive / saturation inside the resonator feedback path.
`Polarity`	Saw / Square	Feedback polarity of the resonator. Saw is positive feedback (softer, vowel-like); Square is negative feedback (hollow, square-wave character).

The Resonator runs post-grain, so its pitch follows the grain material by default. Lower `Mix` values keep the grains in front; crank it toward 1 to turn the plugin into a pure physical-modelled voice.

11.4 Sub Oscillator

A global monophonic sub that tracks the played note and thickens the low end. Set `Level` to 0 to disable.

Control	Range	Description
`Level`	0–1	Output level of the sub. 0 silences it.
`Note`	0–127 (MIDI)	Base pitch of the sub when no MIDI note is playing (Free mode), or the reference pitch in Instrument mode. Displayed as a note name.
`Wave`	Sine / Tri / Saw / Sqr	Oscillator waveform. Sine is pure, Triangle is soft, Saw and Square (both band-limited) add harmonics.
`Octave`	−2 / −1 / 0	Octave offset relative to the played note.

11.5 VCA Envelope

A global ADSR that shapes the amplitude of every played note. In **Instrument** play mode the envelope retriggers on each note-on; in **Free** mode it effectively holds at sustain while the global transport is running.

Control	Range	Default	Description
`A` (Attack)	0.001–3 s	1 ms	Time from note-on to peak level.
`D` (Decay)	0.001–5 s	1 ms	Time from peak down to the sustain level.
`S` (Sustain)	0–1	1.0	Level held while the note is active.
`R` (Release)	0.001–5 s	300 ms	Time to fade to silence after note-off.

11.6 Filter Envelope

A second ADSR that modulates the per-layer filter cutoff. The filter itself — model, cutoff, resonance, drive — lives on the GRAIN tab FILTER column, along with the bipolar `Depth` knob that scales how much of this envelope reaches the cutoff.

Control	Range	Default	Description
`A` (Attack)	0.001–3 s	10 ms	Time from note-on to peak envelope level.
`D` (Decay)	0.001–5 s	100 ms	Time from peak to sustain.
`S` (Sustain)	0–1	1.0	Level held while the note is active.
`R` (Release)	0.001–5 s	100 ms	Time to fade after note-off.

With `Depth` at 0 the envelope has no effect. Push it positive to open the filter on note-on and let it close back toward the static cutoff; push it negative to invert the shape.

11.7 MIDI Control

Grainulator has two play modes, selected in the header settings popup:

- **Free** — the global transport runs grains continuously; MIDI notes are ignored for pitch.
- **Instrument** — note-on / note-off gate the VCA and filter envelopes, and incoming notes transpose the grain pitch relative to the layer's root note.

The MIDI panel in the SYNTH tab exposes three knobs:

Control	Range	Default	Description
`Root 1`	0–127 (MIDI)	60 (C3)	Root note for **Layer 1**. When you play this MIDI note in Instrument mode, Layer 1 plays its source at unshifted pitch.
`Root 2`	0–127 (MIDI)	60 (C3)	Root note for **Layer 2**, independently settable.
`Glide`	0–2 s	0 s	Portamento time between consecutive MIDI notes. Applied to both layers.

Both root-note knobs render their value as a note name (`C3`, `F#4`, …) rather than a raw number, matching the `Note` knobs on the Resonator and Sub panels.

On every MIDI note-on in Instrument mode the engine calculates `note − root` for each layer and applies that offset on top of the `Pitch` knob from the GRAIN tab — so Layer 1 and Layer 2 can transpose against the same incoming note by setting different root values. The Resonator tracks the note offset from Layer 1's root as a single global reference.

`Glide` smooths pitch transitions between notes at the per-layer level; 0 disables portamento, 2 s gives a long, slow slide. It applies equally to both layers.

11.8 See also

- GRAIN Tab — the per-layer filter (model, cutoff, resonance, drive, envelope depth) lives there
- MOD Tab — route LFOs, envelopes, and step sequencers to resonator, sub, and envelope parameters
- Layer System

12. FX Tab

The FX tab holds Grainulator's post-mix effects chain. Everything on this tab is **global** — both layers are already summed by the time the signal arrives, so there is no per-layer FX. Signal flows strictly left to right: **Master Filter → Compressor → Delay → Reverb**.

The master mixing controls — Layer 1 / Layer 2 faders, pans, and Wet/Dry — used to live on FX but were relocated to the GRAIN tab MASTER column in the 2026-04-07 layout update.

12.1 Master Filter

A single post-mix filter stage applied before the compressor. The filter bypasses itself automatically when `Cutoff` is wide open (20 kHz) and `Drive` is 0, so a fully-open setting costs nothing.

Control	Range	Description
`Model`	10 models	Filter topology: Stilson, Microtracker, Krajeski, MusicDSP, Oberheim, Improved, RKSimulation, Hyperion, DaisyLadder, CytomicSVF. Defaults to CytomicSVF. Same bank as the per-layer filter on the GRAIN tab.
`Cutoff`	20 Hz – 20 kHz	Global filter frequency (log-scaled).
`Reso`	0–1	Resonance at cutoff. High values add a singing, whistling character to the full mix.
`Drive`	0–1	Saturation applied alongside the filter. Adds warmth and gentle compression even with cutoff wide open.

12.2 Compressor

A feed-forward log-domain compressor (Giannoulis / Massberg / Reiss 2012) with stereo-linked peak detection, a fixed 6 dB soft knee, and an optional soft-knee limiter. The character is clean and transparent — this is a glue/bus compressor, not a coloration effect. A live gain reduction meter next to the knob grid shows how hard the detector is working. All six continuous knobs read 0–1 on the UI; the ranges in the table are the mapped musical values.

Control	Range	Description
`ON`	on/off	Master bypass for the compressor stage. Click-free. Default **on**.
`Thresh`	−60 to 0 dB	Level above which gain reduction starts. Default ≈ −3 dB (very light touch).
`Ratio`	1:1 to 20:1	Amount of reduction above the threshold. Default ≈ 4:1.
`Atk`	0.1 – 100 ms (log)	How quickly the envelope follower reacts to peaks. Short attacks catch transients; long attacks let them through for punch.
`Rel`	10 – 1000 ms (log)	How quickly gain recovers after a peak.
`Makeup`	0 – +40 dB	Output gain added after compression to compensate for reduction.
`Mix`	0–1 (dry→wet)	Parallel-compression blend. At 1.0 the signal is fully compressed; lower values blend in the uncompressed signal for NY-style parallel compression. Default **1.0** (fully wet).
`LIM`	on/off	Post-compressor soft-knee limiter. Transparent below ±0.9 (linear passthrough), then a smooth `tanh`-shaped rolloff asymptotic at ±1.0 (0 dBFS). No hard clipping, no square-wave fold, no aliased harmonic breakup when the pre-limiter signal is well above 0 dBFS. Default **on** as a safety net.

The compressor only sees the wet signal; the dry path on the GRAIN tab's Dry/Wet mix skips it, so the unprocessed input is never compressed.

12.3 Delay

A stereo tape-style delay with independent L/R times, four routing modes, and two time-change behaviors. Feedback can be pushed above unity for self-oscillating runaways, and the infinite-loop toggle turns the delay into a captured loop.

Control	Range	Description
`TimeL` / `TimeR`	5 – 2730 ms	Delay time per channel (log-scaled).
`Link`	on/off	When on, TimeR follows TimeL. Turn off to set stereo offsets manually. Default **on**.
`Feedback`	0 – 1.25	Amount fed back into the delay line. Values above 1.0 are self-oscillating — use with care.
`Mod Mode`	Tape, Fade	How the delay reacts when its time parameter is changed. **Tape** re-pitches the tail as the read head moves (classic tape warp). **Fade** crossfades between old and new positions without pitch shift.
`Mode`	Dual, PingPong, Single, Cascade	Routing of the two internal delay lines. **Dual:** independent L/R. **PingPong:** L→R→L bouncing. **Single:** mono delay duplicated to both channels. **Cascade:** the two lines are chained in series for longer or more rhythmic echoes.
`Tone`	−1 to +1	Bipolar tone control on the delayed signal. Negative darkens (low-pass), positive brightens (high-pass).
`Sat`	0–1	Saturation drive inside the feedback loop. Adds warmth and iterative grit as feedback regenerates.
`Sync`	on/off	When on, delay times snap to host tempo. Uses `Clock Div` and `Sync Offset`.
`Clock Div`	1/4 – 4/1	Tempo division when `Sync` is on. Available values: 1/4, 1/3, 1/2, 2/3, 3/4, 1/1, 3/2, 2/1, 3/1, 4/1. Default **1/1**.
`Sync Offset`	−5 to +5	Integer offset from the selected clock division (steps in the divisor table) for swing and pushed/pulled feel.
`Loop`	on/off	Infinite-loop mode — the current delay buffer is captured and plays back indefinitely regardless of feedback.
`Mix`	0–1	Delay wet/dry blend for the delay stage itself (independent of the global Dry/Wet on GRAIN).

The two mode parameters are not duplicates. `Mode` is the **routing** between the two internal delay lines (Dual / PingPong / Single / Cascade). `Mod Mode` is the **time-change behavior**: Tape pitches the tail when you turn the time knob, Fade crossfades without pitch shift.

12.4 Reverb

The reverb panel has an **algorithm selector** that switches between three distinct reverb engines. The surrounding knobs (`Size`, `Damp`, `Width`, `Shimmer`, `Tone`, `Mix`, `Freeze`) are shared across all three algorithms — the UI layout never changes — but the underlying interpretation depends on which engine is active.

Algorithm selector

The three engines are named for their character, not their internal algorithm. Default is **Stormy**.

Algorithm	Engine	Character	Best for
`Misty`	Signalsmith Basics	Clean, modern, diffuse. No shimmer, no freeze behavior.	Transparent ambience, pads, realistic-sounding spaces where the reverb should get out of the way.
`Cloudy`	Mutable Rings (Griesinger)	Smooth, classic plate/hall tone with gentle modulation. No shimmer, no freeze behavior.	Warm vintage-style reverb on grain clouds and sustained textures.
`Stormy`	Oliverb (extended Griesinger/Dattorro)	Rich, heavily-modulated, with pitch-shift shimmer in the feedback loop and infinite freeze.	Shimmer pads, ambient tails, evolving textures, long drones. This is the default and the most Grainulator-like.

Common controls

Control	Range	Description
`Size`	0–1	Room size / decay length. All three engines scale decay time from this.
`Damp`	0–1	High-frequency damping in the feedback loop. Higher values make tails darker over time.
`Width`	0–1	Stereo width / diffusion. Maps to diffuser AP coefficients inside each engine.
`Shimmer`	0–1	Pitch-shift amount in the reverb feedback loop. Only audible on **Stormy** — mapped to approximately ±12 semitones around a center-off point. Ignored by Misty and Cloudy.
`Tone`	−100 to +100	Bipolar post-reverb tone shelf. Negative darkens, positive brightens. Common to all engines.
`Mix`	0–1	Reverb wet/dry blend for the reverb stage. Default **0** — reverb is off until you turn it up.
`Freeze`	on/off	Captures the current reverb state and sustains it indefinitely. Only has effect on **Stormy**: input gain, damping, and high-pass are disabled so the tail holds forever. Misty and Cloudy ignore this toggle.

Notes on per-algorithm behavior

- **Stormy** is where the personality lives: nine random modulation LFOs scattered through the topology, a dual-tap pitch shifter per feedback branch, a HP filter to keep long tails clean, and an infinite freeze. Modulation rate and depth scale automatically with `Size`.
- **Cloudy** is the previous default — the Mutable Rings Griesinger reverb. Smooth and musical, without shimmer or freeze.
- **Misty** is the cleanest of the three; best when you want the reverb to feel like a natural space rather than an effect.
- Switching algorithms briefly clears the newly-selected engine's buffers, so you will hear the tail cut on switch.
- `Shimmer` and `Freeze` remain visible in all modes on purpose — they simply do nothing in Misty and Cloudy.

12.5 See also

- GRAIN Tab — master mixing (layer faders, pans, Wet/Dry) and the per-layer filter
- TEXTURE Tab — resonator and sub stages upstream of the FX chain

13. MOD Tab

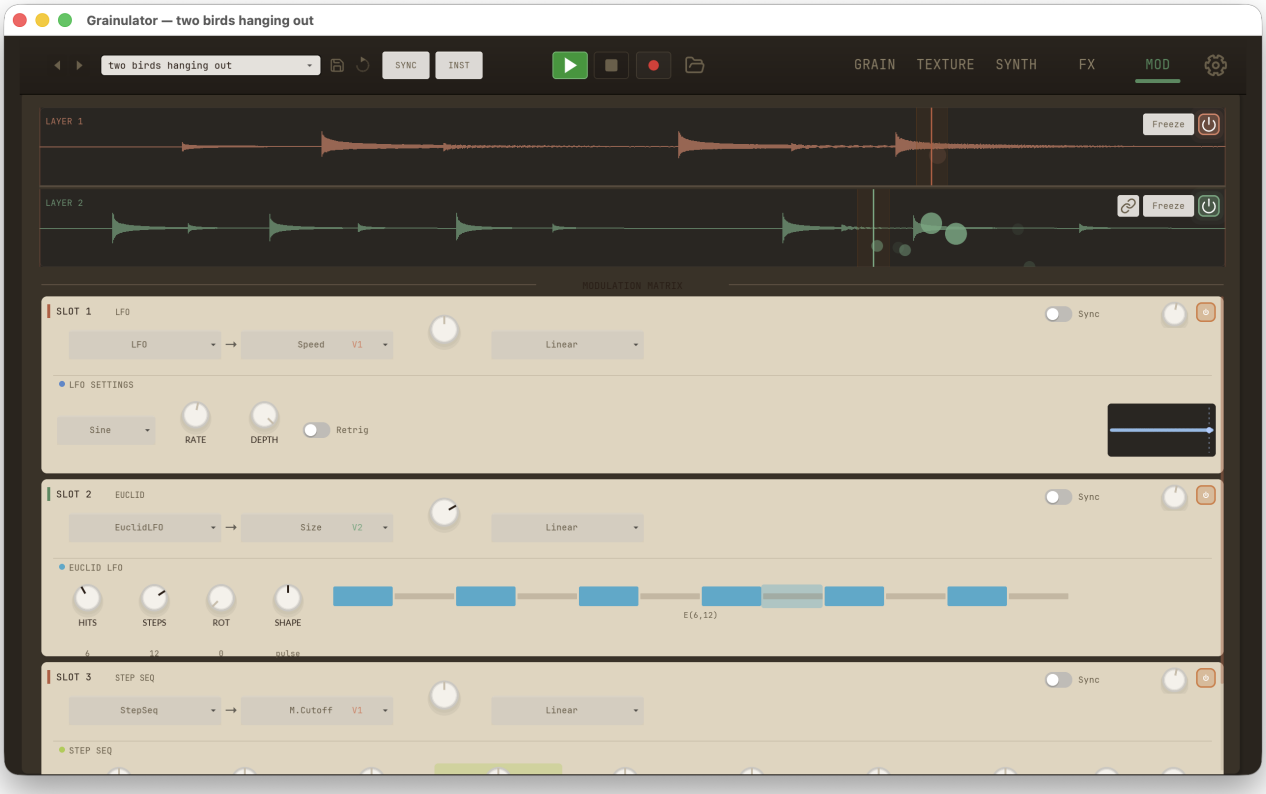

The MOD tab is Grainulator's modulation matrix — five independent **lanes** that each pick a **source**, a **destination**, and a set of per-lane shaping controls. Because the MOD tab is always visible (including in the compact AU layout, where SYNTH / FX / MOD live in a tab strip at the top), every modulation assignment is reachable from any other view without losing your place on the GRAIN, TEXTURE, or FX tabs.

The UI labels each lane `SLOT 1`, `SLOT 2` … through `SLOT 5`; this manual uses "lane" interchangeably with "slot" — they are the same thing.

Each lane carries its own LFO, ADSR envelope, Euclidean generator and 8-step sequencer, so you don't have to reserve a global resource to get movement on a single parameter. Sources cover LFOs, envelopes and two rhythmic generators added on 2026-04-05 — a clock-divided **Euclidean LFO** and a tempo-synced **step sequencer**. Destinations reach **45 active targets** covering the grain core, window, filter, resonator, spectral block, master filter, global effects, cross-voice interaction, and the compressor (plus one **(legacy)** slot retained for v19 preset compatibility — see Destinations below).

Mod lanes always respect the layer model. Each lane has a `Target` that picks which layer(s) it applies to, and the UI uses the **Layer 1 / Layer 2** naming throughout even though the underlying parameters are stored as `Voice 1` / `Voice 2`.

13.1 Overview

Every lane exposes the same set of core controls:

Control	Description
`Source`	What drives the modulation. Per-lane LFO, envelope, Euclidean LFO, or step sequencer.
`Dest`	Which parameter is modulated. One of 45 targets grouped by category.
`Amount`	Bipolar scale from −1 to +1. Negative values invert the source.
`Curve`	Response shaping applied to the source before it hits the destination.
`Target`	Which layer(s) the modulation applies to: both, Layer 1 only, or Layer 2 only.
`Enable`	Lane on/off without having to reset `Source` or `Dest`.

The `Curve` menu offers nine shapes — `Linear`, `Exponential`, `Logarithmic`, `S-Curve`, `Stepped`, `RandomSmooth`, `Threshold`, `Inverted` and `Bipolar`. `Threshold` is particularly handy with the `Spectral Freeze` destination: it converts any bipolar source into a clean gate (1 above zero, 0 below).

13.2 Sources

The `Source` dropdown on each lane lists every available modulation source. The previous Macro / Velocity / Pressure entries are retained in the underlying parameter enum for preset compatibility but are hidden from the UI and produce zero output at runtime — the four live sources are described below.

Per-lane LFO

Classic free-running or tempo-synced LFO. Each lane has its own, so you never run out.

Control	Range	Description
`LFO Shape`	enum	`Sine`, `Triangle`, `Saw`, `Square`, `Random`, `S&H`.
`LFO Rate`	0.01–50 Hz	Free-running rate when the lane's `Sync` toggle is off.
`LFO Depth`	0–1	Output scale applied before `Amount`.
`LFO Retrig`	toggle	When on, the LFO resets to phase 0 on every MIDI note-on (Instrument mode). When off, the LFO runs free.

When the lane's shared `Sync` toggle is on the LFO is driven by the rhythmic clock described below and `LFO Rate` is ignored.

Per-lane envelope

A full ADSR envelope, one per lane. In Instrument mode it is triggered by MIDI note-on and released on note-off; in Free mode it triggers alongside the global play transport.

Control	Range	Description
`Env A`	0.001–10 s	Attack time from 0 to peak.
`Env D`	0.001–10 s	Decay time from peak to sustain.
`Env S`	0–1	Sustain level held while the note is active.
`Env R`	0.001–10 s	Release time after note-off.

Rhythmic sources (added 2026-04-05)

Two phase-driven generators share a single clock with the LFO. The clock can free-run at a Hz rate, or lock to host tempo at a musical division (1/32 through 1/1, straight / dotted / triplet), so all three time-based source types line up with the rest of your project.

Euclidean LFO distributes `Euc Hits` evenly across `Euc Steps` using Bjorklund's algorithm, rotates the pattern by `Euc Rot`, and renders each hit according to `Euc Shape`: at 0 each hit is a gate (square pulse), at 1 it's a ramp that decays linearly to zero over the step, and 0.5 sits between the two.

Step Sequencer plays back an 8-step pattern of bipolar slider values at the lane's clock rate. `Seq Len` truncates the sequence to anywhere from 1 to 8 active steps, and `Glide` slews between adjacent steps over a fraction of the step duration (0 = instant, 1 = full step slew). When the lane's `Dest` is `Pitch` the sliders snap to integer semitones and the hybrid pitch path emits a full ±12 semitone range, bypassing the usual `Curve` and `Amount`.

Every rhythmic source shares the same 18 per-lane parameters:

Control	Range	Description
`Sync`	toggle	Lock this lane's clock to host tempo. Off = use `Free Rate`.
`SyncDiv`	0–15	Division index (1/32 → 1/1, including dotted and triplet variants). Used when `Sync` is on.
`Free Rate`	0.01–50 Hz	Free-running clock rate. Used when `Sync` is off.
`Euc Hits`	0–16	Number of active hits in the Euclidean pattern.
`Euc Steps`	1–16	Pattern length for the Euclidean LFO.
`Euc Rot`	0–15	Pattern rotation. Shifts which step the pattern starts on.
`Euc Shape`	0–1	Hit rendering: 0 = gate, 0.5 = pulse, 1 = ramp.
`Step 0`…`Step 7`	−1…+1	Eight bipolar step values for the sequencer. On pitch destinations they snap to integer semitones.
`Seq Len`	1–8	Active step count. Steps beyond the length are skipped and dimmed in the UI.
`Glide`	0–1	Fraction of step duration spent slewing between step values.
`Reserved`	—	Reserved for future use; has no effect.

One clock "cycle" is one full pass through the source's period — the Euclidean pattern's step count, or the step sequencer's length. A lane synced to 1/4 at 120 BPM therefore produces one full pass every two beats.

13.3 Destinations

Each lane's `Dest` dropdown lists 45 active modulation targets grouped by category, plus one `(legacy)` slot retained for v19 preset compatibility. Per-layer destinations are routed through `Target`; global destinations (master filter, effects, Voice FX, compressor) ignore `Target`.

Group	Destinations
Grain	`Position`, `Speed`, `Pitch`, `Size`, `Density`
Window	`Tilt`, `Curve`, `Morph Amt`, `Morph Rate`
Filter	`Cutoff`, `Resonance`
Resonator	`Chord`, `Brightness`, `Damping`, `Scatter`, `Mix`, `Note`
Amplitude	`Level`, `Pan`, `Spread`
Master	`M.Cutoff`, `M.Reso`
Effects	`Delay Mix`, `Reverb Mix`, `Reverb Size`, `Rev Freeze`
Spectral	`Transpose`, `Tilt`, `Frz Blend`, `Clarity`, `Focus`, `Freeze`, `Formant Mix`, `Partials`, `Peak Lock`, `(legacy)`, `Scan`
Voice FX	`Follow`, `Vocode`, `Duck`
Compressor	`Comp Thresh`, `Comp Ratio`, `Comp Attack`, `Comp Release`, `Comp Makeup`, `Comp Mix`

The **Voice FX** destinations drive the new Voice Interaction panel on the SYNTH tab — see that section for what each target does. The **(legacy)** slot under Spectral is the former `X-Mod` destination, retained so mod-lane routings in pre-v20 presets round-trip; it has no DSP effect in v20+.

13.4 Layer targeting

Each lane's `Target` picks which layer(s) the lane applies to:

Value	Behaviour
`Both`	Apply to both layers.
`Voice 1`	Apply only to Layer 1.
`Voice 2`	Apply only to Layer 2.

The underlying enum labels still read `Voice 1` / `Voice 2`, but the layer they address is Layer 1 / Layer 2 as named everywhere else in the UI. Global destinations (master filter, delay, reverb, compressor) ignore this setting.

13.5 See also

- Layer System
- GRAIN Tab
- TEXTURE Tab

14. Parameter Reference

Grainulator's parameters are documented on the per-tab pages under **Reference → Tabs**. This page is a quick lookup for "which tab is parameter X on?". For ranges, defaults, and descriptive prose, follow the link to the tab page.

14.1 Quick lookup

Area	Parameters	Tab
Grain core	Position, Speed, Pitch, Size, Density, Jitter, Spread, Pan	GRAIN
Per-layer filter	Model, Cutoff, Resonance, Drive, Env Depth	GRAIN
Master mixer	Layer 1/2 Level, Layer 1/2 Pan, Wet/Dry	GRAIN
Global Morph (hero knob)	Macro, Morph LFO Rate, assignable targets	GRAIN
Grain window	Window Shape, Tilt, Curve, Sides, Decay, Window B, Morph, Rate	TEXTURE
Pitch variation	Prob, Mode, Direction, Range	TEXTURE
Dynamics	Prob, Range	TEXTURE
Reverse	Reverse probability	TEXTURE
Scan	Rate, Depth	TEXTURE
Spectral Array Oscillator (per-layer)	Enable, SAO, Freeze, Focus, Transpose, Tilt, FMT mode, Clarity, Partials, Scan, Capture, Clear	SYNTH
Voice Interaction	Enable, Direction, Follow, Vocode, Duck	SYNTH
Resonator	Enable, Polyphony, Chord, Bright, Damp, Scatter, Note, Mix, Distort, Polarity	SYNTH
Sub oscillator	Level, Waveform, Octave, Note	SYNTH
VCA envelope	Attack, Decay, Sustain, Release	SYNTH
Filter envelope	Attack, Decay, Sustain, Release	SYNTH
MIDI	Root 1, Root 2, Glide	SYNTH
Master filter	Model, Cutoff, Resonance, Drive	FX
Compressor	Enable, Threshold, Ratio, Attack, Release, Makeup, Mix, Limiter	FX
Delay	Time L/R, Link, Feedback, Mode, ModMode, Tone, Sat Drive, Sync, Clock Div, Sync Offset, Infinite, Mix	FX
Reverb	Mode (Misty/Cloudy/Stormy), Size, Damping, Width, Modulation, Tone, Mix, Freeze	FX
Mod matrix — per lane	Source, Dest, Amount, Curve, Target, Enable	MOD
Mod matrix — per lane LFO	Rate, Shape, Depth, Retrigger	MOD
Mod matrix — per lane envelope	A, D, S, R	MOD
Rhythmic sources	Sync, Div, Free Rate, Euclid Hits/Steps/Rotation/Shape, 8 Step values, Step Length, Step Glide	MOD

14.2 Layer scope at a glance

Most parameters exist per layer (Layer 1 and Layer 2 can hold different values). The exceptions — parameters that are global regardless of the layer you are editing — are:

- All FX chain parameters (master filter, compressor, delay, reverb)
- The global Morph (Macro knob) and its LFO rate
- Master mix (Layer 1/2 level and pan, Wet/Dry)
- The modulation matrix (though each lane can be targeted at a specific layer via its Target setting)
- **SYNTH tab**: Resonator, Sub Oscillator, VCA envelope, Filter envelope, and Glide are all global even though the UI mounts them under Layer 1. Only Root Note is truly per-layer. The Spectral Array Oscillator controls are genuinely per-layer. The Voice Interaction controls are global (they describe the routing *between* layers, not per-layer state). See Layer System for the full story.

14.3 See also

- Interface — Header Bar
- Interface — Waveform Display
- Layer System

> **Note**
>
> The authoritative ranges, defaults, and units for every parameter live on the per-tab pages linked above. A consolidated auto-generated parameter table is on the post-launch roadmap.

15. iPad install

Grainulator runs on iPadOS as a single universal binary that provides both the standalone app and the AUv3 plug-in. One install gives you both.

15.1 App Store install

Grainulator v2.0 for iPadOS is available on the App Store. Search for **Grainulator** by Plastic Factory and tap install, or open the App Store link from plasticfactory.com/grainulator.

- **Requires iPadOS 15 or later.** Older iPads that can't run iPadOS 15 are not supported.
- **Recommended hardware:** M-series iPad Pro / iPad Air, A17 Pro iPad mini, or 2020 iPad Pro (A12Z) and newer. Older iPads may run Grainulator but CPU headroom for two layers with dense grain clouds is tight.

See performance for detailed hardware guidance.

15.2 First launch

Tap the Grainulator icon on your home screen. The app opens in standalone mode by default — you'll see the compact 5-tab layout (SYNTH / FX / MOD / GRAIN / TEXTURE).

The first time you open the standalone app, iPadOS will ask for microphone permission. Accept it if you plan to record live audio into Grainulator; otherwise you can decline and enable it later under **Settings → Grainulator → Microphone**.

> **Note**
>
> On iPad, Grainulator always uses the compact layout (1280 x 560). There's no tall "desktop" layout — the 5-tab structure is how you move between the synth engine, effects, modulation, grain display, and texture view.

15.3 Using Grainulator as an AUv3 plug-in

Open your AUv3 host, add a new AUv3 instrument slot, and pick **Grainulator** from the list. Each host has its own UI for loading plug-ins — the general pattern is:

1. Create or open a project in your host.
2. Add a new track and choose an AUv3 **instrument** slot on it.
3. Pick **Grainulator** (Plastic Factory) from the instrument list.
4. Tap the plug-in slot to open Grainulator's UI.

Grainulator is a MIDI-controlled instrument, so the host needs to feed it note data — either from a MIDI track, an on-screen keyboard, or an external MIDI controller.

> **Tip**
>
> The standalone app and the AUv3 plug-in share the same preset folder, so presets you save in one are available in the other.

15.4 Sharing audio files into Grainulator

iPadOS sandboxes apps, so audio files need to reach Grainulator through the Files app. Two routes work:

- **Tap the waveform** inside Grainulator to open the iPadOS file browser, then navigate to iCloud Drive, On My iPad, or any configured cloud provider (Dropbox, Google Drive, etc.) and pick a file.
- **Drag-and-drop from Files** while Grainulator is open in Split View or Slide Over. Drop the audio file directly onto the waveform display.

Supported formats: WAV, AIFF, MP3, FLAC, OGG, and M4A.

15.5 Backup and preset sharing

Grainulator presets live inside the app's container. To back them up or share them with another device:

1. Open the Files app.
2. Under **On My iPad**, locate the **Grainulator** folder (created the first time you save a preset).
3. Copy presets to iCloud Drive, AirDrop them, or zip the folder and email it.

Restoring is the reverse — drop the preset files back into the same folder from Files.

16. Touch controls

Grainulator's UI is built around the same gestures on iPad as on desktop — the mouse events map cleanly to single-finger touch. Every gesture listed here is implemented in the current build; nothing is aspirational.

16.1 Knobs

Single tap A single tap on a knob does nothing destructive — it just begins a drag if you keep your finger down. Use it to grab a knob.

Vertical drag Drag your finger up to increase the value, down to decrease. Horizontal movement is ignored. Typical use: set the Pitch, Density, Size, or any other knob. The sensitivity is tuned so a full knob sweep is roughly 200 pixels of vertical travel.

Double tap Double-tap any knob to snap it back to its default value. Useful when you've been sweeping Cutoff or Density exploring sounds and want a clean reset.

> **Fine adjust**
>
> On desktop, holding Shift while dragging gives a 10x-slower fine-adjust. iPad has no Shift key, so fine adjust on touch is not available — land close to the value you want and let the normal drag sensitivity do the rest.

16.2 Buttons, toggles, dropdowns

Single tap Taps activate buttons, flip toggles, and open dropdown menus. Once a dropdown is open, a second tap picks the highlighted option.

16.3 Waveform display

The waveform display (visible on the GRAIN tab and along the bottom of the SYNTH tab) is the only widget with richer gestures.

Tap on an empty layer If a layer has no audio loaded, tapping its region opens the iPadOS file browser so you can pick a WAV/AIFF/MP3/FLAC/OGG/M4A file to load.

Tap near a grain playhead Taps landing near an active playhead grab that playhead. Drag your finger to scrub the grain read position in real time — the grain cloud follows your finger.

Drag across an empty region of the waveform Drags that don't start on a playhead paint a loop region: finger-down sets the start, finger-up sets the end. The grain engine confines reads to this region until you paint a new one.

Drag an audio file onto the waveform From Files in Split View or Slide Over, long-press a supported audio file to pick it up, then drag it onto Grainulator's waveform and release. The file is loaded into the first empty layer (Layer 1 if both are empty, Layer 2 if Layer 1 is already loaded), or into the currently focused layer if both have audio.

> **Drop targeting**
>
> The drag-and-drop handler loads the file into the first empty layer, then falls back to the focused layer. There's no per-half targeting — dropping on the Layer 2 strip won't force-load into Layer 2 if Layer 1 is empty.

16.4 Gestures not implemented

For completeness, these are gestures that exist on some plug-ins but **not** on Grainulator:

- Long-press on a knob does not open a mod-assignment menu. Use the MOD tab to route modulation.
- Two-finger drag is not a separate fine-adjust gesture.
- Pinch-to-zoom on the waveform is not implemented.

17. Performance

Granular synthesis is CPU-hungry. Grainulator is tuned to run on modern iPads, but two layers with dense grain clouds and the resonator enabled will push any iPad hard. This page explains which settings cost the most and how to tune for your device.

17.1 iPad model recommendations

Recommended (comfortable headroom). M-series iPad Pro, M-series iPad Air, A17 Pro iPad mini. These chips handle two layers with high grain density plus resonator plus reverb without drama, even at 256-sample buffers.

Works but tight. A14 through A16 devices — 2020 iPad Air, 2020/2021 iPad Pro, standard iPad (10th gen), 2021 iPad mini. You'll get one full-density layer comfortably; for two layers, expect to dial back grain density or bypass the resonator. Use a 512-sample buffer for live play and bump to 1024 if you're mixing.

Not supported. Pre-A14 devices (anything older than the 2020 iPad Air), or any iPad that can't run iPadOS 15.

17.2 What costs CPU

Rough order of magnitude, from most expensive to cheapest:

1. Grain density. This is the single biggest CPU driver on each layer. Each active grain is a separate read from the buffer plus envelope and pitch-shift math. Doubling the density roughly doubles that layer's CPU cost.

2. Running two layers. Layer 1 and Layer 2 are fully independent grain engines. Running both uses roughly twice the CPU of running one. If you only need a single texture, toggle Layer 2 off from the power button on the waveform display.

3. The resonator. Enabling the resonator adds a fixed, non-trivial CPU cost — around the same as adding another half-layer of grains. If you're thermally throttled or close to the edge, bypass the resonator first.

4. Filter and drive. The moog-style filter and drive stage are cheap. Leave them on.

5. Reverb freeze. Reverb freeze holds the tail indefinitely but doesn't cost more CPU than the normal reverb — it's safe to leave enabled.

17.3 Buffer size guidance

In an AUv3 host you can usually set the audio buffer size in the host's audio settings (128 / 256 / 512 / 1024 samples). Larger buffers lower CPU load but add latency.

- **128 samples:** don't. Too tight for granular on iPad.
- **256 samples:** good for live play on M-series iPads. Expect glitches on older hardware.
- **512 samples:** comfortable on A14-A16 for two layers.
- **1024 samples:** use when mixing / not playing live, for maximum CPU headroom.

The standalone app picks its buffer automatically based on the iPadOS audio session — generally 512 samples at 48 kHz. You can't override this in the standalone app; use an AUv3 host if you need tighter latency.

17.4 Background audio

The standalone Grainulator app registers an audio session and keeps audio running when the app is backgrounded — useful for hand-off to a mixer app or checking something in the Files app while a drone plays. To stop audio, bring Grainulator back to the foreground and press stop, or kill the app from the app switcher.

17.5 Thermal throttling

iPads throttle aggressively when their sustained power draw gets too high, especially the mini and standard iPad. Symptoms of thermal throttling with Grainulator running:

- Audio dropouts and clicks that weren't there a minute ago.
- Sudden CPU usage spikes reported by the host.
- The iPad feeling warm to the touch.

Recovery:

1. Reduce grain density on both layers.
2. Bypass the resonator if it's enabled.
3. Turn off Layer 2 if you can live without it.
4. Let the device cool for about two minutes before pushing it again.

> **Tip**
>
> Running Grainulator from a cool iPad in an AC-conditioned room is very different from running it in direct sun on a patio. If you're playing live outdoors, pre-configure a lower-density preset as your safe fallback.

18. Troubleshooting (iPad)

Quick answers to the most common iPad-specific issues. For general DSP / preset questions, see the desktop troubleshooting page — most answers there apply on iPad too.

"I can't load an audio file"

iPadOS sandboxes apps, so Grainulator can only read files that are explicitly handed to it — usually through the Files app.

First, check where the file actually lives. Open the Files app and confirm the audio file is under **iCloud Drive**, **On My iPad**, or a cloud provider you've enabled (Dropbox, Google Drive, etc.). Files that live only inside another app's private storage (for example, an old voice memo that was never exported) are invisible to Grainulator.

If the file is visible in Files but tapping the Grainulator waveform to browse doesn't show it, try the other route: open Files in Split View alongside Grainulator, long-press the audio file to pick it up, then drag it onto Grainulator's waveform. That drag-and-drop path routes through iPadOS's shared drop region and handles files from more locations than the in-app browser.

Supported extensions are WAV, AIFF, MP3, FLAC, OGG, and M4A.

"No sound from the AUv3 plug-in"

Three things to check, in order.

First, verify the host's audio routing. The Grainulator AUv3 slot has to be on a track that's actually connected to your audio output — in most hosts that means the track isn't muted, its output isn't set to a disconnected bus, and the master output is up. Every host has its own UI for this; check the host manual if you're unsure.

Second, confirm Grainulator is receiving MIDI. Grainulator is a MIDI-controlled instrument, so it only makes noise when notes come in. Play a note from the host's on-screen keyboard or a connected controller. If no grains trigger, the plug-in isn't getting MIDI — check the track's MIDI input routing.

Third, open Grainulator and make sure its own output level isn't muted or turned all the way down on the master section.

"Crackling or dropouts"

This is almost always CPU overload. Check, in order:

- **Buffer size too small.** Bump the host's audio buffer to 512 or 1024 samples. 256 is fine on M-series iPads but too tight on older hardware.
- **Too many layers active.** If you only need one texture, disable Layer 2 with the power button on the waveform.
- **Resonator enabled with short damping.** Short damping times keep the resonator ringing harder, which costs more CPU. Either bypass the resonator or raise the damping.
- **Grain density too high.** Density is the largest single CPU driver. Pull it back on both layers.

If the dropouts only appear after a few minutes of playing, you're probably thermally throttling — see the performance page for recovery steps.

"The app crashed"

Most Grainulator crashes on iPad are memory pressure on older devices.

Before anything else, close other apps from the app switcher to free memory, then restart the host (for AUv3) or Grainulator itself (for standalone). If the crash repeats with only Grainulator running, you've likely found a real bug — please file a report.

To send us the crash log:

1. Open **Settings → Privacy & Security → Analytics & Improvements → Analytics Data**.
2. Scroll to find entries starting with **Grainulator-**.
3. Tap the most recent one, then the share icon, and email it to support@plasticfactory.com along with a brief note about what you were doing when it crashed.

"MIDI isn't working"

If notes from your keyboard or sequencer aren't playing Grainulator:

First, check the MIDI source in the host. The track containing Grainulator needs a MIDI input routed from either an external device, an on-screen keyboard, or another MIDI track. AUv3 hosts vary on how they express this — look for a MIDI-in dropdown on the track or plug-in slot.

Second, confirm Grainulator is in the right play mode. Grainulator has a play mode control that switches between **Instrument** (notes trigger grains at MIDI pitch) and **Free** (grains run continuously regardless of MIDI). If you're in Free mode, incoming MIDI is ignored by design. Switch to Instrument to accept notes.

Third, check the root note. Grainulator pitches grains relative to a root note; if the root is set to something absurd (very high or very low), incoming notes may be pitch-shifted so far they're inaudible. Reset the root note to a sensible value (double-tap the knob for its default) and try again.

19. What's New

A summary of user-visible changes per Grainulator release. For commit-level history, see the GitHub commit log.

19.1 v2.1.0 — April 2026

Cross-voice interaction gets a proper panel, compressor clipping gets fixed, and the MOD + TEXTURE tabs get label polish.

New features

- **Voice Interaction panel** (SYNTH tab) — replaces the previous MASTER · XMOD slot between the two Spectral columns. Three categorically different cross-layer effects — **Follow** (the follower's grain engine reads the driver's live audio), **Vocode** (spectral envelope transfer), and **Duck** (amplitude sidechain, bipolar) — driven by a shared 3-state direction switch (`V1→V2 / ↔ / V2→V1`) and a power toggle that hard-bypasses the whole section. Presets saved with a non-zero X-Mod value are auto-upgraded on load: **Vocode** is seeded from the previous X-Mod amount and direction defaults to **Both**. See SYNTH tab → Voice Interaction panel.
- **Voice FX mod destinations** — **Follow**, **Vocode**, and **Duck** are each routable from any MOD lane under a new **Voice FX** category.

Fixes

- **Compressor limiter no longer distorts when overdriven.** The "brickwall limiter" option used to be a literal hard clip at ±1.0, which folded the waveform into a near-square and produced harsh aliased harmonics when the pre-limiter signal was well above 0 dBFS. It is now a soft-knee limiter — transparent below ±0.9, with a smooth `tanh`-shaped rolloff asymptotic at ±1.0. No breakup on heavily compressed or overdriven material.

UI polish

- **TEXTURE tab** — the **DYN / REV** group labels are spelled out in full as **DYNAMICS** and **REVERSE**. The knob labels in each of those groups now read **Probability** (previously **Prob** and **Reverse**) to match the value they actually control.
- **GRAIN tab — Morph strip** — the Morph LFO rate slider sits directly beneath the shape-wave controls (sin / tri / saw / sqr / s&h) and the visualisation bars moved below. Easier to see cause and effect at a glance.
- **MOD tab — Spectral destination labels** now match the on-screen Spectral panel knob text: **Transpose**, **Clarity**, **Frz Blend**, **Formant Mix**, **Peak Lock**. The former per-voice `X-Mod` destination is renamed **(legacy)** and has no DSP effect in v20+; old mod-lane routings pointing at it still load cleanly.

Preset compatibility

- **Format version bumped to v20.** v19 presets load cleanly — the legacy per-voice X-Mod parameters are retained in the param enum as no-ops, and a one-time migration maps their value onto the new **Vocode** knob so cross-voice coupling you had saved still sounds roughly equivalent on first load.

19.2 v2.0.0 — 2026

A major layout and effects update plus the first public release of the new documentation system.

New features

- **Compressor on the FX tab** — a feed-forward compressor with parallel mix and limiter, sitting between the master filter and the delay in the post-mix chain.
- **Multi-mode reverb** — the reverb now has three selectable algorithms (Misty, Cloudy, Stormy) sharing one unified control surface. Pick the mode that fits the material; common controls behave as expected.
- **Rhythmic mod sources** — clock-divided modulation patterns, including a Euclidean LFO generator and an 8-step bipolar sequencer, each optionally synced to host tempo.
- **Dual waveform display** on the GRAIN tab — see both layers' audio at a glance in the full layout.
- **MIDI root-note glide** — independent root notes for Layer 1 and Layer 2 with a shared glide knob for smooth voice tracking.
- **Layer Link** — Layer 2 can share Layer 1's audio buffer, making it easier to build dense textures from a single source sample.

Layout changes

- **Master panel moved** — the master faders (Layer 1/2 level, pan, wet/dry) relocated from the FX tab to the GRAIN tab's new MASTER column. The FX tab is now dedicated to the post-mix effects chain.
- **Window shape visualization** — a live preview above the TEXTURE tab's shape controls updates as you change tilt, curve, and sides.

Renames and terminology

- **Grainulator** — the project formerly known as *GrainulatorV2* is now shipped simply as Grainulator. Internal parameter IDs still begin with `V1` / `V2` and some code comments reference the old name, but all user-facing copy has been updated.
- **Layer 1 / Layer 2** — what used to be called "Voice 1" and "Voice 2" in the UI is now "Layer 1" and "Layer 2". The terminology change clarifies that a layer is a granular engine instance, not a polyphonic voice in the classical sense.

Documentation

- **First release of the new docs site** — a single source of prose produces the website you are reading, plus two edition PDFs (Desktop and iPad). The old `docs/manual/mac.html` and `docs/manual/ipad.html` files are deprecated; anything they said and the current code disagreed on has been reconciled in favor of the code.

19.3 v1.4 and earlier

Earlier releases are summarised in the project's GitHub Releases page.

20. Credits

Grainulator is made by **Andy Smith** under the **Plastic Factory** label.

20.1 Open-source components

Grainulator stands on the shoulders of several open-source projects:

- **iPlug2** — the cross-format plug-in framework. VST3, AudioUnit, AAX, AUv3, and standalone targets all share a single C++ codebase through iPlug2.
- **Visage** — the GPU-backed graphics and UI library used for every control, panel, and visualization in the editor.
- **CloudSeed** — a classic algorithmic reverb by Valdemar Erlingsson, adapted as one of Grainulator's reverb modes.
- **Mutable Instruments Rings** — the physical-modelling resonator at the heart of the SYNTH tab's Resonator section is a port of the Rings Eurorack module's DSP by Émilie Gillet.
- **Signalsmith Audio** — basics and reverb building blocks used in one of the multi-mode reverb algorithms.

20.2 Documentation

This documentation is built with:

- **MkDocs Material** — the site generator, theme, and search.
- **mkdocs-with-pdf** and **WeasyPrint** — the PDF renderer that produces the Desktop Edition and iPad Edition PDFs.
- **Vercel** — the hosting platform serving `docs.plasticfactory.com`.
- **GitHub Actions** — the continuous-integration pipeline that builds and deploys the site on every push to `main`.

20.3 Thanks

To the beta testers who caught real bugs and to everyone who asked for features that turned out to be better than the ones on the original roadmap. The rhythmic mod sources, the compressor's parallel mix, and the multi-mode reverb all came out of user requests.

If you find a bug, have a feature idea, or want to contribute, file an issue at github.com/azsmith/grainulator-v2.