Anima Manual

What Anima is

Anima is a **physical-modeling synthesizer** for iPad, standalone app and AUv3 instrument. There are no oscillators. Sound happens the way it happens in the physical world: a source of **energy** (breath, a strike) excites an **object** (a bank of tuned, decaying resonances, the modes of a string, a bar, a bell, a pipe), and the result rings out into a **room**.

That's the whole interface: **ENERGY** → **OBJECT** → **ROOM**, left to right, cause to effect.

The OBJECT is a *modal resonator*: 24 coupled resonant modes per voice whose frequencies, loudness, and decay follow physical laws you bend with the knobs. Because mode frequencies are computed directly from the note, no delay lines, no lookup tables, every pitch from MIDI 0 to 127 is exactly in tune, including any microtonal pitch an MPE controller sends. Up to 64 voices of polyphony, and you choose the count.

Anima is the companion instrument to **Hexatone**: plug Hexatone's MPE output into Anima and every microtonal pitch just works, per-note, with no configuration.

The control language

Every control is one **glyph**: a bold shape whose form *is* the value. There are no knobs and no numbers at rest; the panel is a picture of the current sound.

- **Touch a control** and it wakes: the plate lights, the glyph begins to move, and the label shows its current value. It only starts *changing* once your finger travels a little, so a press alone lets you see what a control is set to without nudging it.

- **Drag vertically** anywhere on a control's plate to change it (up = more). The plate brightens, the label shows the value, and a small bubble floats above your finger.
- **Double-tap** resets a control to its default.
- **Loading a preset morphs the whole surface**: every glyph glides to its new shape over a quarter second, so you *see* the surface become the new sound. (The audio switches instantly; only the drawing eases.)
- **Undo / redo** (the two arrows in the top bar): one drag, preset load, or reset is one step. A change you regret, or a preset loaded over an unsaved sound, is one tap back.
- **Panic** (the stop circle at the top right): the first tap stops all voices immediately, release tails and held notes included; the reverb keeps ringing. A second tap clears the reverb too.
- **Glide** (the small time beside Voices): portamento. Off, each note starts exactly where it should; raise it and every new note slides in from the last one you played, taking that long to arrive. Drag it sideways (right = longer), or tap it: a selector opens with an off chip and a slider for the time. It works for the ribbon, the keys, and anything playing Anima over MIDI, and MPE bends ride on top untouched. Glide is part of the sound: your saves carry it, and factory sounds, designed glide-free, load with it off.
- **Voices** (the numbers beside Glide): the small figure is how many voices are sounding right now (ringing tails included), the large one is your polyphony ceiling, 1 to 64. Tap it and a selector opens right below: arrows step the ceiling, MONO drops to 1, MAX opens all 64. Dragging the number sideways still works (right = more). At 1 the synth is a last-note mono instrument; lowering it mid-chord releases the notes above the new limit. It's a setup choice, so presets never change it.

Color = family

In the **Colored** palette (the default), each control is inked in the earthy pastel of its **family** so related controls rhyme in hue, and each hue keeps to a section, so a column reads as one temperature:

Color	Family	Controls
sand	breath / energy	Force · Noise · Flutter · Inertia · Envelope · Curve
clay	touch / impact	Strike · Contact
ochre	brightness (ENERGY)	Edge · Peak · Edge Env
heather	the object's body	Rigidity · Harmonics · Density · Twin · Bloom · Ring Time · Tone
dusty denim	the room's time	Pre-delay · Size · Regen · Decay
bone	gain	Level · Mix
sage	the room itself	Blur · Diffusion · Dark · Width · Color

Reading by color: sand is *energy over time* (breath and its envelope), heather is *the object* (what rings and how its partials sit), sage and denim are *the room*. The rows are ordered so the families sit in blocks: each row reads as a gradient of what the controls do, left to right.

Prefer one ink? **SETUP ▸ LOOK ▸ Monochrome** drops the families and inks every glyph in a single colour you pick from a wheel; the labels stay neutral for legibility. **Field** there is Night (soft dark) or OLED (true black). Everything applies instantly.

Orientation

- **Landscape:** three full-width rows (ENERGY / OBJECT / ROOM).
- **Portrait:** three columns read left to right, controls flowing downward, same controls, same order.

ENERGY: how the object is excited

The energy source is a continuous pressure/noise stream shaped by an envelope, plus an optional strike. This is "how you play it": blow at it, press it, hit it.

Envelope: the wide strip. The solid silhouette is the ADSR; drag its **three handles**:

- **Attack apex** (top-left dot): drag sideways for attack time (1 ms – 5 s).
- **Knee** (the corner where decay meets the sustain shelf): a free 2-D point, sideways sets decay (1 ms – 5 s), up/down sets sustain level.
- **Release foot** (bottom-right dot): sideways for release (1 ms – 10 s). The nearest handle to your finger wins, anywhere on the strip. Double-tap near a handle to reset that segment.

Curve: the envelope's shape, spanning the whole family: **exponential** at the bottom (fast start, slow tail, the snappy, "acoustic" feel), **straight lines at noon**, **logarithmic** at the top (slow start, fast finish, deliberate swells, notes that gather themselves before speaking). Each segment keeps its exact time at any curve; only the path bends.

Force (0–2), how hard the energy drives the object. Above 1 it's genuinely harder playing, not just louder: more energy into every mode.

Noise: what the breath is made of. 0 = pure steady pressure (its onset edge gives a soft press/thump). Turning it up doesn't just add hiss, it *widens the turbulence*: first a dark rumble that gives the press its weight, then a shaded, wooden breath through the middle, opening into full white air at the top. The knob fades the noise's color, not its amount, so every position is a usable texture. Most acoustic sounds live in between.

Flutter: breath instability, and a *dual-purpose* travel: one drag sets both **how fast** and **how deep** the pressure wanders. The rate climbs steadily across the knob (≈0.5 Hz at the bottom to ≈8 Hz at the top) while the depth **arches**: nothing at 0, the full wobble at *noon*, then trading depth for speed until only a fast, shallow shimmer remains at the top. So: left half = slow, deepening sway; middle = the breath at its most alive; right half = nervous tremor. Because the wander also reaches the breath that feeds Inertia, a note near the sing threshold flickers in

and out of singing, like a player at the edge of their air. At 0 the pressure is dead steady, exactly as before. (The strike is exempt, an impact has no breath.)

Strike: a contact pulse at note-on, and the knob is *the mallet itself*: low = a soft felt press (a long, dark thud), high = a hard wooden crack (short and bright). It happens *immediately*, even under a slow pad attack, the impact is physical, the breath is musical. Velocity rides on top, making hits both louder and brighter (a shorter pulse has a wider spectrum). At 0 there is no strike at all, pure breath.

Edge: the brightness of the contact: a filter on the energy stream, from soft felt to glassy hiss. Displayed in Hz.

Peak: resonance at the Edge frequency: a formant-like emphasis, the "shape of the mouth." It focuses rather than boosts: raising it narrows the breath toward the Edge frequency without making the notes that land there jump out louder than their neighbours.

Edge Env: how much the *loudness envelope also drives Edge*. Physically: blowing harder doesn't just make breath louder, it makes it **brighter**: more hiss, more edge. At 0, brightness is frozen wherever Edge sits, no matter how the note evolves. Turned up, the ADSR sweeps Edge upward (as much as four octaves at full): the attack **blooms bright**, the sustain settles darker, the release dims to a whisper, all in step with the volume contour you drew. To hear it isolated: slow attack, Edge low, Edge Env high, each note starts as a dark pressure and *opens* into breath.

Off-panel (host-automatable): Velocity sensitivity (1 = fully velocity-scaled, 0 = flat).

OBJECT: what is being played

Ten controls describing a physical object. Read the row as a sentence: *how long it rings, what it's made of, how its partials sit, how many it has, where it's touched, how it's voiced, whether it's one object or two, how it blooms, how it breathes, how loud it is.*

Ring Time: how long the object rings after being excited: 50 ms (dead wood) to 25 s (a great bell). This is the object's own decay, independent of the room.

Rigidity: the material, expressed as *which partials die first*. Low = soft wood: the highs vanish quickly, leaving a warm thump. High = glass/metal: everything rings out evenly, cold and bright. (High notes always die faster than low ones, like real materials.)

Harmonics: the spacing of the partials, morphing from a perfect harmonic series (strings, pipes) through stretched piano-wire tension into fully inharmonic bell/bar spacing. The fundamental never moves, only the overtones migrate.

Density: how many partials the object owns: from ~4 (clean tuned percussion, kalimba, woodblock) to the full 24 (rich, complex). Loudness stays constant as modes come and go.

Contact: *where* the object is touched. At the edge, every partial speaks. Sliding toward the center progressively silences the even partials, at dead center the object is hollow, clarinet-like. (The glyph tips like a see-saw: touch at the edge throws it off balance.)

Tone: the spectral tilt of the partials, –9 to +6 dB per octave: dark felt to glassy presence. This is the object's voicing; Edge is the *exciter's* brightness, they're the same orange family for a reason.

Twin: a second, detuned copy of the object (1.5–21 cents). Small amounts breathe and chorus; larger amounts beat audibly, bell doubles, honky-tonk, unison width.

Bloom: nonlinear coupling that pushes energy *upward* through the partials over time: strike low, and the top swells in a moment later, the way gongs and cymbals bloom. Physically bounded, it can never run away, at any setting.

Inertia: regenerative breath, on an exponential knob. The lower two-thirds is a tonal region: the feedback adds body and resonance, and the note still always decays. **Around 70 percent it crosses into self-oscillation**: the object sings like an overblown reed or a bowed plate as long as breath is applied (it always dies on release). That threshold sits in the same place for every material, tuning, and register. **The top third drives the singing voice into saturation**, squaring it up into a genuine overblown scream, thick with harmonics, by full. As you climb, the level is compensated automatically, so the sound gets more intense rather than just louder, it never blows up, and it can't run away.

Level: the object's voice, pre-room. (The glyph is Anima's eye, opening.)

Off-panel: Tune (±24 st), Fine (±100 c), Drift, a slow, living detune walk of the upper partials (the fundamental never drifts), like tape or old organs.

ROOM: where it happens

Not a synthesis stage, the ready-made space that makes everything sound finished. A 16-line feedback-delay-network reverb fronted by a colored echo.

Mix: the dry/wet balance between the direct sound and the ROOM. The crossfade is loudness-compensated (equal-power, wet leg level-matched), so moving Mix changes the balance, not the overall level. **Size**: how big the space is: the reverb's dimensions and the echo's own delay time (20 to 380 ms), one knob for both. Sweep it while repeats ring and they bend like tape. **Decay**: the halo tail's length, 0.2 to ~16 s (the echo's repeats belong to Regen). **Diffusion**: how quickly reflections smear into wash.

Regen: how much of each echo repeat survives into the next. 0 is a single slapback; the middle is a polite trail; past ~90 percent the loop crosses into **real feedback territory**: the repeats build instead of dying, saturating a little on every pass until they park at a singing ceiling, and they stay there until you back the knob off (with Color up, each pass also darkens, the way bucket brigades do). At very short Sizes the loop is quietly reined in so it can never turn into a screaming comb.

Dark: high-frequency damping, down to a genuinely drastic felt-curtain at full (it darkens the echoes too, not just the tail).

Width: stereo width of the room **and of the dry voices**: notes place themselves across the image by pitch, scaled by Width. It also ping-pongs the echo's repeats across the image. At 0 everything is honestly mono.

Blur: the morph between the two worlds: 0 is the bare delay (discrete colored repeats), full blurs it into the reverb's halo (pure wash). The delay's time and

repeats stay put (Size and Regen own those); Blur chooses how much of each world you hear.

Pre-delay: the gap before the room answers, 0–200 ms. At 0 the room speaks instantly; at higher values the dry note stands apart from its space.

Color: the bucket-brigade haze. Saturation plus a woody high-frequency loss that accumulates with every repeat, and past the middle the echo becomes a real bucket brigade: its clock wanders, so repeats drift and shimmer slightly in pitch, and the loop resamples itself coarsely enough that bright material folds down into a soft metallic sizzle. That aliasing is deliberate; it is the haze. The tank keeps only the gentler grit, and Freeze is immune (a hold holds). The echoes age as they repeat.

Freeze: off the panel for now (its cell made room for Regen). It lives on as a host-automatable parameter in AUv3, where it still holds the room's contents indefinitely and frosts the ROOM glyphs while held.

Presets

Tap the **preset name** (top center) for the library; the **gear** (top right) opens SETUP (look, manual, about):

- **Init**: the designed default patch, always first in the list. A fresh install opens on it.
- **Save**: name the current sound, tap Save. Yours appear above the factory banks, with delete. If the name already exists, Anima asks whether to overwrite that save or keep both.
- **Copy patch**: puts the *current* sound, saved or not, on the clipboard as text. Universal Clipboard carries it between your devices; paste into Notes/Messages to collect keepers (each copy includes a human-readable summary of every setting, and the text loads back losslessly). **If you find a sound you love, Copy patch first, tweak second.**

- **Factory banks**: the built-in sounds, grouped in named banks (Stock first; artist packs join as their own sections).
- **Export presets**: shares all your saved sounds as one JSON file (AirDrop, Files), backup, transfer, and the pipeline for turning personal presets into factory ones.
- Presets capture the **sound** (ENERGY/OBJECT/ROOM, plus Glide and the ribbon's snap mode), never your tuning, MPE, or Freeze state. Factory sounds are glide-free by design, so loading one turns Glide off; factory presets leave your snap mode alone.
- Standalone and AUv3 share one preset library.
- In hosts, the stock bank also appears as native **factory presets**, and full state saves/restores with the session.

The ribbon (standalone)

A continuous pitch surface along the bottom:

- **X = pitch**, continuously, slide for true glide. Fret lines with C labels for orientation.
- **Multi-touch**: every finger is its own voice with independent glide.
- **Slide up or down** and the slide axis moves, relative to where you landed: the same axis an MPE controller's slide (CC74) sends, under your finger. Where it goes is yours to route: tap **SLIDE ROUTING** in the ribbon's header and send it to any mix of the ENERGY and OBJECT controls, each with its own depth (zero is off; Tone at 0.50 is the factory setting). The envelope is not routable, and Strike fires at note-on, before a slide exists, so it is absent too. Controls with a routing wear a small dot on the panel and move as you slide; the ribbon stays playable under the panel, so set depths by ear.
- **Snap** is a three-way switch, all three settings on the strip: Off, Snap, Loose (in a narrow window it collapses to one chip that cycles on tap). **Snap** quantizes touches to semitones, and the fret you are snapped to stays lit while your finger is down (and under held notes while Hold is on), so the semitone is never

a mystery. **Loose** is the third way: you land on the fret, moving tracks your finger exactly (a true gliss), and when you stop, the pitch settles onto the nearest fret in a breath; the light appears as the note arrives. Lift near a fret and the settle completes, so held notes land in tune. Your choice is remembered, and saves with your presets.

- The ribbon owns the bottom of the screen: the first swipe up from the bottom edge plays the ribbon, and only a second, deliberate swipe reaches the iOS app switcher. Guard screen edges (SETUP) extends the same rule to every edge, Control Center included. One gesture stays out of any app's reach: the multi-finger app switch. To lock the screen completely while performing, turn on Guided Access once in Settings ▸ Accessibility, then triple-click the side button in Anima.
- **Hold** (hand icon) latches notes: build chords one finger at a time; tap a held note to release just it; slide-then-lift leaves it where you let go.
- **Octave arrows** shift the range.
- **Keys** (piano icon) swaps the ribbon for a three octave piano keyboard in the same frame: real, wide black keys, the same octave arrows and Hold, and sliding across keys still glides chromatically. Tap the icon again for the ribbon; your choice is remembered.

MIDI & MPE

- **Ordinary MIDI just works**: channel-agnostic note on/off, velocity, pitch bend, all-notes-off. Sequencers and keyboards need zero setup.
- **Program Change selects a factory preset**: program number N loads the Nth factory preset, counting across the banks in order. Send it from a host's MIDI track or a controller's patch buttons to step the banks hands-free (numbers past the last preset are ignored).
- **MPE is automatic**: when a controller announces MPE (Hexatone, LinnStrument, Seaboard, ROLI Blocks…), Anima switches itself, per-note pitch bend becomes true per-note microtonal pitch. In SETUP's MIDI section, MPE mode Auto/On/Off covers controllers that never announce themselves.

- **What each MPE axis does**: **strike** (how hard you land) is velocity; **pressing into a held note is breath**, per note (sustained patches swell and soften under the finger; a patch with Inertia near its threshold ignites and sings as you press; percussive patches have little breath for pressure to act on, which is physics, not a fault); **the slide axis (CC74) goes where you route it** (the ribbon's SLIDE ROUTING panel: any mix of the ENERGY and OBJECT controls), per note, while its neighbours keep theirs.
- **Depths are setup, not presets**: the Pressure slider lives in SETUP's MIDI section, and the slide axis's depths live in the ribbon's SLIDE ROUTING panel; together they set how strongly your controller's touch reaches the sound. They are calibration for your hands and hardware, so a preset load never changes them and they survive relaunches (0 turns an axis off). All of them are also Anima parameters a host can automate.
- **Glide and the bend range**: sliding along an MPE surface glides that note's pitch continuously. Anima follows the range the controller announces (RPN 0), defaulting to the MPE standard of 48. A slide that comes out as tiny vibrato means the controller and synth disagree about the range: set the controller to 48 in its editor, or match Anima to it (SETUP ▸ MIDI ▸ Bend range). On a plain (non-MPE) keyboard the pitch wheel uses the standard 2 semitones, and follows whatever range the keyboard announces.
- **Bluetooth MIDI**: SETUP ▸ MIDI ▸ Bluetooth MIDI devices opens the system pairing screen; pairing is system-wide, and the same section lists everything currently connected. Inside a host you will not see the input list or the pairing button: the host owns MIDI routing there, so route your controller to Anima in the host and play.
- **Sustain pedal (CC64)** holds released notes until the pedal lifts, exactly as a piano player expects. Panic and all-notes-off always win over the pedal.
- **Hexatone → Anima**: route Hexatone's MIDI to Anima in AUM (or via the virtual "Anima" destination) and every hex is in tune. Nothing to configure.
- Standalone receives CoreMIDI (USB, Bluetooth, virtual, network) and keeps sounding **in the background** while you work in another app.

AUv3

Anima loads as an `aumu` instrument in AUM, GarageBand, Logic, Cubasis, etc. All parameters are host-automatable, and the controls on screen follow host automation and LFOs: a modulated control moves and shows its live value under the glyph, the same as under a finger. Factory presets appear natively; sessions restore exactly. The plugin window adapts to any shape the host gives it (row layout when wide, column layout when tall, scaling below that). By default the plugin shows the surface alone (your host owns the keyboards); pull the small handle at the bottom edge, or tap it, and the ribbon slides up so the plugin plays exactly like the standalone, keys mode, Snap and all. Anima first asks the host for a taller window so nothing is covered (AUM says yes); hosts that refuse get the strip floating over the bottom edge instead. The choice is remembered.

Sound-design recipes

Breathy flute: Noise high (~0.8), Inertia just past 0.5, Ring Time short-ish, slowish attack (~100 ms), Edge mid, Contact off-center. The note *sings* while held and dies with the breath. Add **Flutter** (~0.5) and the singing wavers like a real player's air.

Struck bell: Sustain 0, short decay, **Strike up**, Harmonics high (bell spacing), Ring Time long, Rigidity high (metal), a little Twin, Contact near the edge.

Blooming gong: As the bell, but a low note, **Bloom 0.6+**, some Drift, Ring Time maxed. Strike it and listen to the top arrive late.

Sparse tuned percussion: **Density low** (~0.2), Rigidity low (wood), Ring Time short, Strike up. Kalimba/woodblock territory.

Hollow woodwind: **Contact at center** (evens null out), Noise high, moderate Inertia: instant clarinet family.

Overblow: Any patch: push **Inertia past 0.5** while holding a note, then keep going, the upper half of the knob drives the sing into a squared-up scream. Bright/metal

patches are the filthiest; the fundamental always leads. With **Flutter** up, a note parked *at* 0.5 flickers across the threshold by itself.

Frozen choir (AUv3): Play a chord into a big room, automate **Freeze** on from your host, play a melody over the held wash.

Dub machine: **Blur at 0** (pure delay), **Regen high**, **Color high**: repeats that darken and saturate as they age. **Size** is your delay time; sweep it mid-trail for tape bends. Push Regen past 90 percent and the loop self-oscillates: it builds to a singing ceiling and holds until you pull back. Ride the knob.

Living drone: Long Ring, Inertia ~0.5, host-automate **Drift** up: the partials start breathing like tape.

Knobs that talk to each other

Anima's controls are physical laws sharing one object, so some pairs interact in ways no single knob's description tells you. The surface knows these relationships: while you hold a knob, the plates of its companions warm faintly, a quiet hint at what to reach for next. This chapter is the why. None of this is required; all of it is playable.

- **Force and Bloom.** Bloom listens to the *level* of the fundamental, squared. So Force doesn't just make a blooming patch louder, it makes the overtones arrive sooner and harder: blow gently and the top barely stirs, blow hard and it blossoms.
- **The envelope and Inertia.** The envelope isn't only loudness, it is the breath that feeds the sing. A note ignites when the envelope stands high enough, so **Sustain decides whether the sing survives past the attack**: set it low and a note can catch fire on the attack peak and collapse into its own sustain. If a singing patch seems to die mid-note, look at Sustain before you blame Inertia.
- **Flutter and Inertia.** Near the sing threshold the breath wander makes a note flicker in and out of singing by itself, like a player at the very edge of their air. The most alive setting in the synth.

- **Noise and Inertia.** Noise at 0 with Inertia past the bridge is the pure sung tone: nothing drives the object but its own feedback. The cleanest flute in the box.
- **Strike and Attack.** The strike fires at the instant of contact even under a slow swell: impact first, breath blooming after. A slow pad with a small Strike reads as "touched, then breathed on."
- **Edge Env and Curve.** With a slow, logarithmic attack (Curve high) and Edge Env up, a note gathers itself in the dark and *opens* into brightness exactly as it swells.
- **Twin and Ring Time.** The twin's beating needs time to be heard: at short rings it reads as chorus, at long rings the beat outlives the breath and the release keeps shimmering on its own.
- **Contact and Harmonics.** Contact at center hollows the even partials (the clarinet trick); add Harmonics stretch on top and the odd stack bends bell-wards. Two knobs, four instrument families.
- **Rigidity and Ring Time.** Ring Time says how long, Rigidity says *who*: soft materials spend the ring on the lows (the highs die at once), metal and glass let everything ring. The same Ring Time reads completely different across the Rigidity travel.
- **Regen, Size and Blur.** Size is the delay's time and Regen its feedback, so the dub tricks live between them: self-oscillation is earned at longer Sizes, and Blur decides whether the repeats stay repeats or melt into the wash.

Inside the engine: how Anima actually makes sound

This section is for the musician who wants to know what's really in the box. Nothing here is required to play Anima, but knowing what the knobs are made of changes how you reach for them.

Why there are no oscillators

Almost every synthesizer starts from an oscillator: a circuit that already *is* a sound, a buzzing sawtooth, a hollow square, which filters then carve away at. That's a picture of sound as a substance you sculpt.

The acoustic world doesn't work that way. A flute is not a tone generator; it's a column of air and a stream of breath. A bell is not a waveform; it's a piece of bronze and a moment of contact. In the physical world, **sound is what happens when shapeless energy meets a resonant object in a space**. The energy has no pitch. The object has no sound. Only the encounter sings.

Anima is built on that sentence, literally. The three panel sections are its three nouns:

- **ENERGY** makes something *shapeless*, pressure and noise, the acoustic equivalent of white light.
- **OBJECT** is a simulated resonant body that turns that energy into pitch and timbre by *remembering* some frequencies and forgetting the rest.
- **ROOM** is where it happens.

When you turn an Anima knob you are almost never "adding an effect." You are changing a physical law (the material, the geometry, the touch), and the sound follows the way it would in the world.

ENERGY: breath is white light

The exciter produces a mix of two things: **steady pressure** (a constant push, like air against a reed) and **noise**, the turbulence. At the top of the Noise knob the turbulence is white (every frequency at once, in equal measure, which is exactly why it's the perfect raw material: whatever the object wants to resonate at is already in there). Below that, the knob doesn't thin the noise out, it **narrows its band downward**: the density stays the same, only the width shrinks, from open air through a shaded, wooden breath down to a sub-bass rumble riding the press. A dark breath still drives the low resonances at full strength; what it loses is only the top.

A subtlety worth knowing: the object can't hear *steady* pressure at all. Like your ears, it only responds to change. So the pressure component speaks at its **edges**: the moment the envelope rises (the press) and falls (the release thump). That's why low-Noise settings sound like touch and high-Noise settings sound like breath.

Before the energy reaches the object, it passes through a filter, the **mouth**. Edge is that filter's brightness, Peak its formant-like emphasis, Edge Env its link to effort (blow harder = brighter, mechanically coupled to the envelope). The Envelope itself is not "a volume control over time"; it *is* the act of playing: how you press, hold, and let go of the air. Curve bends its path from exponential snap up into logarithmic hesitation without changing its timing.

Two things bypass parts of this chain on purpose, because physics says so. **Strike** is a contact pulse whose brightness comes from its *duration*: a hard, short impact contains high frequencies not because a filter opened but because narrowness in time is wideness in frequency. It skips the envelope (an impact happens at the moment of contact, even under a slow swell) and skips the mouth filter (its width *is* its spectrum). And **Flutter** wobbles the pressure itself (no real player's breath is a laboratory constant), which also means it wobbles everything downstream that feeds on breath, Inertia included.

OBJECT: a thing made of 24 resonances

Here is the heart of it. Physics has a beautiful theorem: **any vibrating object (a string, a bar, a bell, a drumhead) can be completely described as a set of independent resonances** ("modes"), each with three numbers: a frequency, a loudness, and a decay time. The wood, the shape, the size: all of it collapses into that list. Change the list, change the object.

Anima's OBJECT is exactly that list, made playable: **24 resonators per voice** (48 when Twin engages a second, detuned copy). Each one is a tiny mathematical flywheel, a point rotating in a circle at its mode's frequency, losing a little energy each turn, kicked by the incoming breath. When you play a note, nothing is "looked up": the 24 frequencies are *computed from the pitch you asked for*, every note, continuously. This is why every MIDI note from 0 to 127 is exactly in tune, why an MPE bend from Hexatone lands on the precise microtonal frequency it asked for, and why glides never step or zipper: the flywheels re-tune smoothly mid-spin.

Each OBJECT knob edits one column of the list:

- **Ring Time** scales all the decay times, how long the object remembers.

- **Rigidity** tilts decay *against frequency*. This is what "material" actually is, acoustically: in soft wood, high modes die almost instantly (the thud); in glass and metal they ring nearly as long as the fundamental (the shimmer). And in every real material, higher notes die faster than low ones. Anima keeps that law, referenced to middle C.
- **Harmonics** bends the *frequency* column. At zero the modes sit at 1×, 2×, 3×⋯ the fundamental: the harmonic series, the signature of strings and air columns. Turning up first adds the gentle stretch of stiff piano wire, then breaks into the clustered, inharmonic spacing of bars and bells. The fundamental is mathematically pinned (it cannot move), so the note never drifts while its overtones migrate.
- **Density** decides how many of the 24 exist at all, fading them in and out with loudness conserved: four modes is a kalimba, twenty-four is a cathedral bell.
- **Contact** applies the touch-point law. Where you excite an object determines which modes can speak: a mode with a motionless point (a node) under your finger stays silent. Touch a string at its middle and every even harmonic dies. That's the hollow, clarinet-like sound at Contact's center. The comb follows the real formula, with one act of mercy: the fundamental is exempt, so no touch point can ever silence the note itself.
- **Tone** tilts the loudness column: the object's voicing, dark felt to glassy presence. (Edge brightens the *breath*; Tone brightens the *body*. Same orange family, different physics.)
- **Twin** runs the entire second bank a few cents away: two objects, never quite identical, beating against each other the way two real bells always do.
- **Bloom** is the one deliberately *nonlinear* coupling: it listens to the fundamental's swing and pushes energy **upward**, at twice the frequency, harder the louder it swings. That's the gong law: strike low, and the top arrives a breath later. Because the coupling is square-law, it does nothing at small amplitudes: it cannot self-ignite, only respond.
- **Inertia** closes the loop that turns an object into an *instrument*: the object's own vibration is fed back into itself, scaled by your breath. Below the threshold, losses win: the note always decays. When gain meets loss the object crosses into **self-oscillation**: it sings for as long as breath is applied, the way a reed speaks or a bowed plate hums. The loop is calibrated so that threshold sits in

the same place *for every material, tuning, and register*. That consistency is engineering, not accident. The knob is exponential, so that crossing lands around 70 percent of its travel, leaving the lower range as a tonal region. Past it, the loop is driven into saturation: the singing waveform squares up and screams, with the output level compensated as you climb so the top of the knob adds intensity, not decibels, and never blows up.

- **Drift**, off-panel, lets each upper mode's frequency take a slow private walk (the object breathing like tape), the fundamental exempt so the tuning stands still.

One philosophical consequence worth noticing: because the OBJECT is a list of physical properties rather than a waveform, **a preset in Anima is not a sound. It's a thing.** It behaves differently at every pitch, every velocity, every touch, the way a real marimba bar does.

ROOM: a space, not an effect

The ROOM is sixteen delay lines cross-connected so that every line continuously feeds every other (through a Hadamard matrix, the mathematically even-handed way to mix sixteen channels). That tangle of mutual echo is what a room *is*: reflections reflecting reflections until individual paths blur into a wash. Size sets the lines' lengths (and the echo's own time: one knob for how big the space is), Decay the halo's regeneration, Regen the echo's, Dark and Color how each round trip loses brightness and gains character: repeats that age, like tape. Blur slides one knob between the two worlds, from **discrete colored echo** (the reflections still distinguishable) to **pure halo** (fully blurred: the knob is named for exactly this). Freeze (a host-automation parameter now) closes every loss in the loop: the room's current contents circulate untouched, a held breath of space. And Width is honest: at zero, even the voices' own stereo placement collapses; the instrument is truly mono.

The whole voice, end to end

As many voices run in parallel as you set, up to 64, each one the full story: breath → mouth → object → its own place in the stereo image (voices pan by pitch, low left of high, scaled by Width). The sum self-levels (a chord of ten is balanced against a solo note by measured power, not guesswork), and the output stage

protects your ears twice: a transparent safety limiter rides *sustained* extremes down (normal playing never touches it), and a soft clip rounds any transient that still tries to slam. Every control you touch is applied on the audio thread between samples, so nothing ever clicks; every frequency is recomputed live, so nothing is ever out of tune.

Nothing in Anima replays a stored sound. Every note is a small physical event, energy meeting an object in a space, computed while you listen. That's the reason patches feel like *instruments*: what you're adjusting isn't a recording's dressing, it's the laws of a little world.

The mathematics of a note

One level deeper: not the numbers this time (those follow below) but the ideas behind them. Everything Anima computes at audio rate grows out of one small mathematical object, and this section is that object explained. Nothing here is needed to play.

A resonance is a rotating point. Take a point on a plane, some distance from the center. Each sample, do two things to it: rotate it around the center by a small fixed angle, and pull it inward by a whisker. Now watch only its height while this repeats tens of thousands of times a second: the height traces a sine wave that dies away exponentially. A pure decaying tone, which is exactly what one resonance of a bell or a string is. The angle per step sets the frequency; the inward pull sets the decay. That is the entire audio-rate machinery of the OBJECT: a mode is this point, its two coordinates are the mode's whole memory, the breath arrives as little kicks that knock the point outward, and one coordinate is what you hear. Twenty-four of them per voice, each with its own angle and its own pull, all kicked by the same breath, their heights summed: that sum is the object's voice.

Why rotation, and not something cleverer. There are many textbook ways to build a digital resonator, and most can misbehave when re-tuned while they ring: a coefficient change at the wrong moment adds energy, and added energy inside a feedback structure is how instruments explode. The rotating point cannot do this, for a reason visible without algebra: a rotation never changes a point's distance from the center, and the inward pull only ever shrinks it. Whatever sequence of angles you dial in, mid-glide, mid-bend, mid-drag, the point can only travel inward.

Stability is not tuned into Anima; it is a property of circles. The whole instrument leans on that one decision, because Anima re-tunes constantly: every glide, every MPE bend, every Harmonics drag re-angles all twenty-four rotors while they are still spinning.

How close to forever a note lives. Ring time is measured as T60, the time to fall by 60 decibels, a thousandth of the starting amplitude. Exponential decay means losing the same fraction every step, so the per-sample pull follows from the ring time directly, and the numbers are worth feeling once: a two-second ring at 48,000 samples per second means keeping about 99.9928 percent of the radius each step. All of music's sustain lives in the hair's breadth between "almost one" and "one". At exactly one a note is eternal; a fraction above, an explosion. Anima spends its life steering inside that gap: Ring Time and Rigidity set the distance from the line, and Inertia, below, walks up to it on purpose.

What noise does to a resonator. The breath is noise: every frequency at once, none of them for long. A resonator listens narrowly, accumulating only the noise that lands near its own frequency, a random walk that builds over its memory span. The mathematics says the built-up level grows with the square root of the ring time: modes that remember longer collect more, and left alone that would quietly turn Ring Time into a loudness knob. Anima divides the growth back out, so ring time reads as duration. The same idea explains the Noise knob's color law: below the white zone the band narrows at constant density, so a mode still inside the band is driven exactly as hard as before, and the fade to darkness you hear is modes dropping outside the band, not the breath getting quieter.

Brightness is shortness. A pulse cannot be both brief and narrow-band: a contact lasting some time t necessarily contains frequencies up to roughly 1/t. That trade is a theorem, not a design choice, and Strike is built directly on it. The knob shortens the contact pulse and the spectrum widens in exact proportion; no filter opens, time itself does the equalizing.

Why the singing settles instead of exploding. Inertia closes a loop: the bank's own vibration, scaled by your breath, is fed back in as more excitation. The oscillation condition is the classical one: a note sustains itself when the energy returned per cycle matches the energy lost, arriving in step with the vibration. Two mathematical facts shape the loop. First, phase: of the rotor's two

coordinates, only one moves in step with the drive at resonance; the other, the one you listen to, runs a quarter turn behind, and a loop closed on it can never satisfy the condition, only damp. Anima feeds back the in-step coordinate. Second, the ceiling: the loop passes through a soft saturation, and a soft ceiling's effective gain falls as the swing grows. As the singing builds, the loop's gain slides down until it equals exactly one, and there the amplitude parks: a stable, breath-sized sing instead of a runaway. The loop is re-calibrated live, solving the gain condition backward from the knob, which is why the threshold sits at the same point of travel for every material, register, and sample rate. The ROOM's Regen uses the identical mathematics past unity: feedback above one, parked by saturation.

Doubling by squaring. Bloom squares the fundamental's motion, and squaring a sine hides an old identity: it yields a tone at twice the frequency, plus a constant that is removed. Energy climbs an octave without any oscillator being asked to exist. And because the square of a small number is negligible, the coupling has no small-signal gain at all: it is mathematically incapable of igniting from silence and can only answer motion that already exists.

Sixteen echoes that lose nothing in the mixing. The ROOM's tank folds its sixteen delay lines into one another every pass through a Hadamard matrix, an orthonormal transform: the sixteen-dimensional equivalent of a pure rotation. Like the rotor's rotation it conserves energy exactly, so however densely the echoes interleave, the mixing itself neither adds nor removes a single decibel. All loss is placed in one deliberate spot, the per-line gains, which is why the tail's decay is designed rather than emergent. And the line lengths are prime numbers scaled by Size: no two share a common divisor, so repeats never gang up on a shared period and the tail stays a wash instead of a flutter.

The ear hears squares. One habit runs through everything: loudness follows energy, and energy is amplitude squared. So every crossfade and pan in Anima keeps the squares summing to one: Twin's two banks, each voice's stereo placement, a chord's self-leveling against a solo note (a sum of N voices is scaled by one over the square root of N). Whenever two things trade places in Anima and the loudness holds still, that is why.

None of this is a simplified imitation of some finer truth. The rotating point is the same mathematics a physicist writes down for a real vibrating bar, and the loop condition is the one that governs a bowed string or a flute's air column. Anima just solves them live, twenty-four modes per voice, up to sixty-four voices, tens of thousands of times a second, while you listen.

The specifics, for the curious

The numbers behind the nouns. None of this is needed to play, but it's the honest account of what happens between samples.

How the resonance is made. Each of the 24 modes is a *coupled-form complex resonator*, picture a point rotating on a circle at the mode's frequency, its radius shrinking a hair each turn so energy bleeds away. Two state variables (an in-phase cosine and a quadrature sine) are advanced every sample by a rotation; the breath arrives as a kick, and the cosine component is the mode's output. This form stays stable *while it is being re-tuned mid-note*, which is what lets glides and MPE bends slide the whole bank smoothly instead of zippering. The three columns of the mode list are computed live from the note you play:

- **Frequencies:** `f_k = f0 · k^(1+w) · √((1 + B·k²)/(1 + B))`. At rest (Harmonics = 0) that reduces to `f0 · k`, the pure harmonic series. Harmonics raises the stretch exponent `w` and the stiffness term `B` together, first bending the partials sharp like tensioned piano wire, then clustering them into inharmonic bell and bar spacing. The `√((1+B·k²)/(1+B))` normalization pins `f_1 ≡ f0` exactly, so the fundamental never moves however far the overtones migrate.
- **Amplitudes:** `A_k = k^(−β) · |sin(π·k·x_e)|`. The `k^(−β)` term is the spectral tilt (Tone), dark and felted at high β, glassy and present at low. The `|sin(π·k·x_e)|` term is the touch-point comb (Contact): a mode whose node falls under the excitation point `x_e` is silenced, which is why Contact at center hollows out the even partials into a clarinet. The fundamental is exempt from the comb, so no touch point can ever mute the note itself.
- **Decays:** `T60_k = T60 · (f_k / 261.6 Hz)^(−σ)`, clamped to 0.005–40 s. `σ` is the material slope (Rigidity): large σ makes high modes die faster than low (soft wood, the thud), σ near zero lets them all ring together (glass, metal, the

shimmer). It's referenced to middle C, and the upward multiplier is capped at 4× so a bass note on soft material can't ring for a runaway minute.

Coefficients are recomputed on a 16-sample grid and linearly interpolated between hops, so nothing steps; pitch is smoothed by a 3 ms one-pole with a 600 semitone/second slew ceiling. Modes crossing 0.40–0.45 × the sample rate fade out and are culled, so even MIDI 127 is a clean, in-tune note with no aliased partials. A fixed output calibration (∝ √(1/samplerate)) holds loudness constant across sample rates.

The exciter and its noise. The energy source is a blend, set by Noise, of **steady DC pressure** and **noise**. From about two thirds of the knob upward the noise is white: a flat spectrum, equal energy at every frequency, so the object always finds something to resonate with. Below that seam the noise leg runs through a one-pole low-pass whose cutoff falls exponentially from ~20 kHz down to ~80 Hz at the bottom, at *constant passband density*: the band narrows but what remains drives its resonances exactly as hard, so the knob is a color fade (air, shaded breath, rumble, press), not a volume fade, and it behaves identically at every sample rate. The last tenth of the travel fades the rumble out so 0 is the pure press. The exciter is high-passed at ~16 Hz, which is precisely why steady pressure is inaudible to the object except at the envelope's moving edges (press and release). It then passes a two-pole low-pass "mouth": **Edge** is its cutoff (mapped as a musical note), **Peak** its resonance/Q, and **Edge Env** couples the cutoff to the envelope for up to four octaves of "blow harder, brighter." Peak's gain is normalized above its default: raising it narrows and colors the band but the resonant peak itself never exceeds the default's level, so a note whose partial rides the Edge frequency focuses instead of jumping out (the wolf-tone fix). **Flutter** is breath instability made from white noise through two variance-normalized one-pole low-passes (at the flutter rate and 3.2× it), giving a band-limited wander whose rate ramps 0.5 → 8 Hz across the knob while its depth arches (nothing at the ends, ±50% at noon); it modulates both the breath level and the self-oscillation drive, clamped so it can never push the loop past threshold. **Strike** is a raised-cosine contact pulse whose length runs from ~7.8 ms (a soft felt press) down to ~0.8 ms (a hard crack), hardness scaled by velocity, and it deliberately bypasses both the envelope and the mouth filter, because an impact's brightness *is* its shortness in time. The ADSR is phase-based (each

segment lasts its set time at any curve), and Curve is bipolar: exponential snap at one end, a true time-mirrored logarithmic hesitation at the other, linear at noon.

The resonator bank. 24 modes per voice, doubled to 48 when Twin runs a second bank detuned a few cents and equal-power mixed, the slow beating of two real bells. **Density** continuously fades the live mode count from 4 to 24 with a smoothstep and re-levels loudness as modes come and go, so a four-mode kalimba and a full cathedral bell sit at the same level; the fundamental is never faded. The regenerative breath (**Inertia**) feeds the bank's cosine output through a tanh back into the injection, re-calibrated every control hop so the small-signal loop gain *at the fundamental* is exactly $L = 4 \cdot \text{inertia}^2$. Self-oscillation therefore begins at inertia = 0.5 for every material, register, and sample rate; the knob is quadratic, so that lands near 70 percent of its travel and the lower range stays a tonal region. It must tap the in-phase cosine component; the quadrature tap can never satisfy the oscillation condition and only damps. Each mode rides a per-mode feedback cap so no overtone can out-sing the fundamental, and the threshold accounts for both Twin banks (a Lorentzian weighting of the detuned copy) so a full Twin doesn't secretly double the loop. Past the threshold, an RMS-invariant waveshaper *outside* the loop squares the tone up toward a scream, and a gain trim keyed to how far past you are compensates the level, so the drive intensifies without blowing up (filth, not decibels). **Bloom** is a separate square-law coupling ($\tanh(3 \cdot u_1)^2$, DC-blocked at 8 Hz) that pushes fundamental energy up an octave; being quadratic it has zero small-signal gain, so it can respond but never self-ignite. **Drift** (off-panel) gives each upper mode a slow private detune walk on golden-ratio-spaced LFOs (0.13–0.74 Hz), depth growing toward the top of the spectrum, the fundamental exempt.

The room. ROOM is a 16-line feedback delay network. The delay lengths are distinct primes scaled by **Size** (no two lines share a period, so the tail stays dense and free of flutter echo), mixed every pass by a 16×16 orthonormal **Hadamard** matrix, the maximally even way to scatter sixteen channels into one another, and lossless by construction, so stability comes only from the per-line feedback gains (all < 1, set by **Decay**). Each line carries a one-pole damping filter (**Dark**, with a 550 Hz floor) for frequency-dependent decay, plus a slow, staggered delay-length modulation (0.08–0.30 Hz) that keeps the tail from ringing metallic. In front of the tank sits a colored stereo **pre-echo**: a feedback delay whose repeats accumulate **Color**: a BBD/tape-style tanh saturation and progressive high-frequency loss, the

woody aging you hear building on each repeat. The echo is a real delay: **Size** sets its time (20 to 380 ms, the same knob that scales the tank's primes, fractional and per-sample smoothed, so a sweep glides like tape), and **Regen** sets its feedback: identical to the proven 0.92-linear law through most of the travel, then climbing through unity to about 1.15 at full. Past unity a soft saturation stage in the loop (engaged only above 0.9 feedback, bypassed while frozen) absorbs the excess each pass, so the self-oscillation builds and parks at a bounded ceiling instead of running away. A three-step guard eases the ceiling back down at short times (0.75 below ~30 ms, 0.92 by ~80 ms, self-oscillation earned only above ~140 ms) so tiny times can never scream as a comb. Decay no longer touches the echo; the two tails are independent. **Blur** is a pure output crossfade from the raw colored echo toward the fully diffuse halo: delay and reverb are one structure at two ends of a single knob. **Pre-delay** (0–200 ms, fractional and per-sample smoothed) feeds the tank *directly*, plus a 0.35 bleed of the echo and a 0.30 early tap of the diffused input, so the room can answer within a sample when you ask it to. **Freeze** mutes the input, drives both loops to ~unity (0.99999), bypasses the damping, and cuts the echo-into-tank bleed, so the held wash neither decays nor slowly pumps; it simply hangs. **Width** decorrelates the sixteen outputs across the stereo field and, at zero, collapses everything (including the voices' own pitch-panning) to honest mono.