# Sandbender Manual

Audio Damage, Inc.

Release 1.0




9 November 2026

© 2026 Audio Damage, Inc. All rights reserved.

The information in this document is subject to change without notice and does not represent a commitment on the part of Audio Damage, Inc. The software described by this document is subject to a License Agreement and may not be copied to other media except as specifically allowed in the License Agreement. No part of this publication may be copied, reproduced, or otherwise transmitted or recorded, for any purpose, without prior written permission by Audio Damage, Inc.

# Credits

## Software Design and Construction, Documentation

Chris Randall
Adam Schabtach

## Field Testing

Jens Ågren
Matt Dowdney
Duncan Foster
Anders Göran
Michael Grommon
Chris Hahn
Steve Hamman
David Henkel
Matthew Lyon
Peter Nyboer
Justin Pearson
Konstantin Voinov
Kent Williams

## Made Possible By

Elle
Tracie

# License Agreement

BY INSTALLING THE SOFTWARE, YOU ARE CONSENTING TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS AGREEMENT, THEN RETURN THE PRODUCT TO THE PLACE OF PURCHASE FOR A FULL REFUND.

Single User License Grant: Audio Damage, Inc. (“Audio Damage”) and its suppliers grant to Customer (“Customer”) a nonexclusive and nontransferable license to use the Audio Damage software (“Software”) in object code form solely on computers owned or leased by Customer.

Customer may make one (1) archival copy of the Software provided Customer affixes to such copy all copyright, confidentiality, and proprietary notices that appear on the original.

EXCEPT AS EXPRESSLY AUTHORIZED ABOVE, CUSTOMER SHALL NOT: COPY, IN WHOLE OR IN PART, SOFTWARE OR DOCUMENTATION; MODIFY THE SOFTWARE; REVERSE COMPILE OR REVERSE ASSEMBLE ALL OR ANY PORTION OF THE SOFTWARE; OR RENT, LEASE, DISTRIBUTE, SELL, OR CREATE DERIVATIVE WORKS OF THE SOFTWARE.

Customer agrees that aspects of the licensed materials, including the specific design and structure of individual programs, constitute trade secrets and/or copyrighted material of Audio Damage. Customer agrees not to disclose, provide, or otherwise make available such trade secrets or copyrighted material in any form to any third party without the prior written consent of Audio Damage. Customer agrees to implement reasonable security measures to protect such trade secrets and copyrighted material. Title to Software and documentation shall remain solely with Audio Damage.

LIMITED WARRANTY. Audio Damage warrants that for a period of ninety (90) days from the date of shipment from Audio Damage: (i) the media on which the Software is furnished will be free of defects in materials and workmanship under normal use; and (ii) the Software substantially conforms to its published specifications. Except for the foregoing, the Software is provided AS IS. This limited warranty extends only to Customer as the original licensee. Customer’s exclusive remedy and the entire liability of Audio Damage and its suppliers under this limited warranty will be, at Audio Damage or its service center’s option, repair, replacement, or refund of the Software if reported (or, upon request, returned) to the party supplying the Software to Customer. In no event does Audio Damage warrant that the Software is error free or that Customer will be able to operate the Software without problems or interruptions.

This warranty does not apply if the software (a) has been altered, except by Audio Damage, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Audio Damage, (c) has been subjected to abnormal physical or electrical stress, misuse, negligence, or accident, or (d) is used in ultrahazardous activities.

DISCLAIMER. EXCEPT AS SPECIFIED IN THIS WARRANTY, ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE, ARE HEREBY EXCLUDED TO THE EXTENT ALLOWED BY APPLICABLE LAW.

IN NO EVENT WILL AUDIO DAMAGE OR ITS SUPPLIERS BE LIABLE FOR ANY LOST REVENUE, PROFIT, OR DATA, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGES HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE EVEN IF AUDIO DAMAGE OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In no event shall Audio Damage’s or its suppliers’ liability to Customer, whether in contract, tort (including negligence), or otherwise, exceed the price paid by Customer. The foregoing limitations shall apply even if the above-stated warranty fails of its essential purpose. SOME STATES DO NOT ALLOW LIMITATION OR EXCLUSION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES.

The above warranty DOES NOT apply to any beta software, any software made available for testing or demonstration purposes, any temporary software modules or any software for which Audio Damage does not receive a license fee. All such software products are provided AS IS without any warranty whatsoever.

This License is effective until terminated. Customer may terminate this License at any time by destroying all copies of Software including any documentation. This License will terminate immediately without notice from Audio Damage if Customer fails to comply with any provision of this License. Upon termination, Customer must destroy all copies of Software.

Software, including technical data, is subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licenses to export, re-export, or import Software.

This License shall be governed by and construed in accordance with the laws of the State of Colorado, United States of America, as if performed wholly within the state and without giving effect to the principles of conflict of law. If any portion hereof is found to be void or unenforceable, the remaining provisions of this License shall remain in full force and effect. This License constitutes the entire License between the parties with respect to the use of the Software.

## System Requirements

The following table summarizes the operating system requirements and formats provided for Sandbender. Sandbender is a 64-bit plugin.

| Operating System | Minimum Version | Formats |
|---|---|---|
| macOS | 10.15 | AAX, AudioUnit, CLAP, VST3; Intel and Apple Silicon |
| Windows | 10.0 version 1607 | AAX, CLAP, VST3 |
| Ubuntu | 22.0 | CLAP, LV2, VST3 |
| iOS (separate purchase) | iOS 16 | AUv3 |

## Demonstration Version

We encourage you to download and try the demonstration version of Sandbender before purchasing it. The demo version of Sandbender is the same as the regular version, but has the following limitations:

- Presets cannot be saved, nor can parameter values or other settings. This includes the information usually stored by your host DAW. If you save a DAW session with an instance of the demo version of Sandbender, the plugin will revert to its default state when you reload the session.
- Sandbender will cease to emit audio altogether 20 minutes after you add it to your DAW session. You can remove it and add it again, but it will revert to its default state.

If you purchase Sandbender after using the demonstration version, simply run the installer provided to you after your purchase to replace the demo version with the full version.

## Introduction

One of the definitions for *synthesis* in the Merriam-Webster dictionary is "the combining of often diverse conceptions into a coherent whole … also: the complex so formed."[1] Sandbender is both a synthesizer and a synthesis, combining elements from several traditions of synthesizer design in an instrument that is both familiar and novel.

### Sound Generation

We started with a four-operator FM core, but we didn't hew to the usual architecture of Yamaha's instruments. Instead we based our design on the RYK ALGO[2], a Eurorack module that takes the four-operator idea in a different direction. Each of its four operators has a single warp control that bends the oscillator shape along a continuous path from triangle through sine to pulse, so an operator is not a sine wave with a fixed timbre but a shape you dial in. Operators are wired together by an *algorithm*, as you would expect. There are ten predefined algorithms but you can connect the operators together in many other ways as well. Additionally, each operator can be both a modulator (which changes the timbre of one or more other operators) and a carrier (which generates the signal you hear), contrary to Yamaha's tradition of operators performing only one of these two functions.

Next in the signal path is a wave shaper with three shaping curves: two West Coast-style wavefolders, one symmetric and one not, and a soft saturator. The wave shaper's single control dials in anything from subtle harmonic enhancement to piercingly bright transformation of the sound coming from the four operators.

Yes, we said "piercingly bright", so next up is a fat and flexible analog-modeled filter. Sporting numerous modes and response curves, as well as an input-drive control of its own, this filter can tame the runaway high frequencies that FM and wavefolding are so capable of spewing forth, bringing the sound back to something complex and changing but warm. Or, should the need arise, you can crank up the filter's resonance and drive it hard, adding its raw, analog howl to the digital intricacy of the FM operators. West Coast wavefolding and FM[3] meet East Coast subtractive filtering and the results are the best of both.

Finally, the output stage is a model of the Buchla 292[4] low-pass gate, the vactrol-based circuit that gives West Coast synthesis its particular way of becoming silent. A low-pass gate is not a VCA and it is not a filter; it is both, unavoidably and at the same time, because the photoresistor it is built around gets quieter and darker together. Hit it and it rings down like something struck rather

[1] "Synthesis." Merriam-Webster.com Dictionary, Merriam-Webster, https://www.merriam-webster.com/dictionary/synthesis. Accessed 20 Oct. 2026.

[2] Product name is a trademark of RYK Modular.

[3] It's not often pointed out that FM synthesis, in the DX7 tradition, is just as rooted in the West Coast as Buchla's early work. John Chowning pioneered FM synthesis while at Stanford University. The company Buchla and Associates was based in Berkeley, California; Stanford and Berkeley are about 30 miles apart as the seagull flies.

[4] Product name is a trademark of Buchla U.S.A.

than something deliberately silenced. Sandbender gives you the gate on its own, the gate combined with its low-pass, or a plain clean VCA when you would rather have the note you asked for.

## Modulation

Modulation is just as essential to synthesizing interesting sounds as the sound generation itself, and Sandbender provides a good-sized assortment of modulators: two envelope generators, two LFOs, two function generators, a sixteen-step sequencer and four macro knobs. Any of these can be routed to very nearly any control in the instrument, including the controls of the other modulators. The function generators and the sequencer have input points exactly as a module would, so for example the sequencer's clock can come from a function generator's end-of-rise signal, or a function generator can be triggered by a sequencer gate. A plethora of sources derived from MIDI messages, including a full MPE implementation, put Sandbender under your fingers via your favorite controller. The patches you build become closer to a small modular system than to a modulation matrix.

## Effects

We decided that Sandbender deserved a special complement of effects, so we took half a dozen of our popular effects plugins and built them in. At your disposal are:

- ADverb, our algorithmic reverb capable of subtle spatial enhancement, huge ambient washes, and everything in between
- Dubstation, our much-loved delay, modeled on bucket-brigade delay processors of the last century
- Evil Otto, our berzerk implementation of aggressive multi-band dynamics processing, essential for getting attention in a loud mix
- Kombinat, our multi-band tool of sonic destruction
- Phase Two, our dual phase shifter which started life as a meticulous model of a rare analog pedal
- Quatromod, our four-pack of silky and fluid modulation effects

These are not approximations or "inspired by" versions. They are the same code, running inside Sandbender, and they can be arranged in whatever order you desire.

## Breeding New Sounds

Creating new synthesizer sounds from scratch can be a daunting and sometimes frustrating process. At the same time, using only the presets provided with an instrument can feel a little stale and derivative. Sandbender provides an intelligent tool for generating new patches without requiring detailed tweaking of its numerous controls. A Genetics page enables you to explore new sounds via randomized but directed evolution. Starting with a sound you like (or don't), you explore new sounds related to, but different from, the original. Sounds evolve along different genetic lineages, guiding you into different sonic territories you might never reach via your own programming.

## Performance

The same page also becomes an instrument in its own right. The PERFORM surface gives you two nodes to drag around, each one moving two controls of your choosing at once, and records what you do with them as a loop of up to eight bars that runs in time with your song. Loops are saved with the preset. With a mouse you can play one node at a time; on a touchscreen, both.

## If You Don't Like Reading Manuals

We get it. Sandbender is a large instrument with a correspondingly large manual. You do not have to read it in order, and you certainly do not have to read it before you play, and honestly we might be a little hurt if you didn't just dive in and start making noise. Load a preset, turn the four LEVEL knobs in the operator section, and you will learn more in ten minutes than this document can tell you in an hour. Come back when you want to understand details like what the WARP knob is actually doing, or why the low-pass gate sounds like that, or how to get the sequencer to clock a function generator.

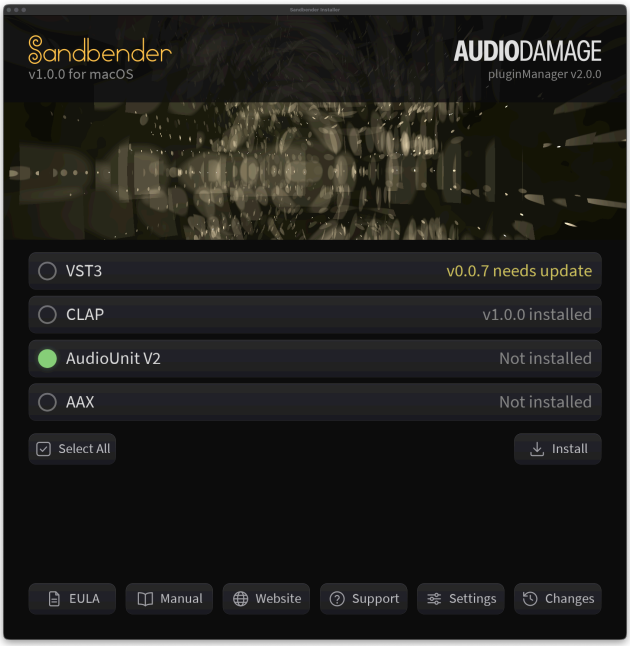
## Historical Anecdote

Audio Damage is Adam and Chris, and as it happens, we both have West Coast roots of our own: we spent much of our formative years in Oregon. We've since made our homes in other parts of the country, but now and then we still miss the Pacific and its sandy, rocky shore.

We hope you enjoy using Sandbender as much as we enjoyed building it.

## Installation

Sandbender uses our custom plugin manager application for installation. Launch it as usual on your operating system of choice and you'll be presented with a window like this:




At the right of the window, beneath the name of the plugin, you'll see the version number of the software carried by the installer. This is distinct from the version of the plugin manager itself, which is shown in the upper right and usually not of much interest. (Note that the version numbers in this screenshot are for demonstration purposes only and may not reflect the actual version numbers of the software at the time you read this manual.)

The installer lists the plugin formats which can be installed: AAX, AU, CLAP, LV2 and/or VST3, depending on the operating system. For each format there's a button on the left to select the format and text on the right describing whether or not the plugin is installed. If the plugin is installed the text will describe whether it's up to date. In the above screenshot, AAX and AudioUnit are not installed, VST3 is installed but older than the version contained in the installer, CLAP is installed and up to date, and AudioUnit is selected for installation.

Click the round button on the left of the format name to select the format for installation. Clicking the SELECT ALL button selects all available formats for installation.

No changes to your system's storage device take place until you click the INSTALL button near the lower-right corner of the window. Click that button and you'll receive visual confirmation that the format(s) you've selected have been installed. On Windows and Linux, if you hold down the SHIFT key on your keyboard, the INSTALL button's label switches to UNINSTALL, and clicking it will remove the selected formats from your system[5]. Once you've installed and/or removed the formats you need, simply close the application in the usual manner for your operating system. You're done. There is no license code or other authorization necessary; we'd rather assume we can trust you than burden you with an onerous DRM system.

You'll find some handy buttons at the bottom of the window, all of which are pretty self-explanatory:

- EULA: presents the End-User License Agreement for our products. By clicking the INSTALL button you're implicitly agreeing to these terms, but we expect that you'll find them reasonable should you take the time to read them.
- MANUAL: opens the current version of this user manual, in PDF form, in your web browser.
- WEBSITE: opens the product's web page in your browser.
- SUPPORT: displays information for contacting us, either via our Discord presence or through email.
- SETTINGS: produces a panel in which you can choose the installation locations for each plugin format. Note that these locations are nominally stipulated by the publishers of the formats and changing them might cause problems for some hosts.
- CHANGES: displays a summary of what we've added, changed or improved since the last version of the product. This button isn't present on the first release of the product.

[5] Blame Apple, not us, for the lack of this feature on macOS. On macOS just manually delete the plugin(s) from your plugin folder(s).

# Operation

Sandbender runs either as a polyphonic instrument with eight voices or a mono synth with a single voice. Each voice has four oscillators working as FM operators, a sub oscillator, a noise source, a filter and an output stage. The modulation sources run within the voices; MIDI sources are applied either to all voices or to individual voices depending on whether or not you're using MPE. The effects are applied globally after the voices are added together.




## Signal Flow

Within a voice, the path is short and runs one way:

The four oscillators are evaluated in order, each one adding whatever signals the algorithm sends it to its own phase, resulting in frequency modulation in the parlance of our time. The operators marked as audible in the algorithm are panned across the stereo field by the SPREAD control and summed onto a bus. That sum goes through the wavefolder, and only then are the sub oscillator and the noise source added to the bus, so neither of them is ever folded. The sum passes through the filter, then through the output stage, which is where the amplitude envelope is applied and where the voice gets its shape in time. The output stage is either a plain VCA or a low-pass gate, selected with a switch.

Two things about the signal path are worth knowing before you start turning knobs. The first is that the FOLD control acts only on the audible sum, never on the signal an operator sends to another operator. An operator that is both heard and modulating feeds the folder with one copy of itself and the next operator with another, unfolded. This is deliberate: folding the modulation path would change every sideband in the algorithm rather than the timbre of the thing you can actually hear, and you would lose the ability to make a sound dirtier without also making it different.

The second is that the filter sits *before* the output stage. With the output stage set to one of the low-pass gate modes, you therefore have two filters in series, one you control and one that moves on its own every time a note is struck. That is not a mistake, but it does explain why an already-dark patch can vanish entirely when you switch the output stage to LPG Both.

The eight voices are summed, the MASTER and LEVEL controls are applied, and the result runs through the effects rack. The effects run in whatever order you have dragged them into. Last of all, after the effects, comes a switchable output limiter to protect your speakers and ears from unpleasant surprises.

Sandbender also applies automatic gain adjustment when playing chords. Eight notes played together are a good deal louder than one, and on most synthesizers a big chord hits everything after it much harder than the single notes you set the patch up with. Sandbender evens this out: whenever more than one voice is sounding, it turns the sum of the voices down until the chord is as loud as its loudest note, judging each voice by its envelope and velocity at that moment. Two equal notes come down 3dB, four come down 6dB, and a single note is never touched. The level comes down quickly enough to catch the attack of a struck chord, and recovers slowly so that letting go of one note doesn't make the ones you're still holding swell. This gain mechanism is always on in POLY mode. In MONO and RETRIG, where only one voice ever sounds, it has nothing to do.

## The Window

Clicking the icons at the top of the window selects one of five pages: Genetics, Synthesizer, Effects, Settings and Info. The Synthesizer page is the instrument proper and is where you will probably spend most of your time. Effects holds the effects rack. Genetics breeds variations on the current sound, and also holds the PERFORM surface, for playing the sound by hand. Settings holds a number of global preferences which are not saved within presets. Info displays the version of the plugin and its credits, with links to Sandbender's page on our website, our Discord server, and email support. The support link starts an email to us with your version of Sandbender, your host and your computer already filled in, which helps us help you when something goes wrong. On a computer, resting the pointer on almost anything in the window brings up a tooltip saying what it is and what it does; if they get in your way, turn them off on the Settings page.

Across the top of every page is the preset display, which doubles as a readout: touch or drag any control and the display shows that control's name and its value until you let go. Click the preset name to open the preset browser. The arrows on either side of the name step backwards and forwards through every preset, traversing folders as needed, in the order the browser lists them but skipping favorites. Step past the last preset and you wrap around to the first, and the other way round. The star just left of the forward arrow marks the current preset as a favorite. The preset browser has a chapter of its own after this one.

In the standalone version of Sandbender, the top bar has one more switch in the form of a keyboard icon. It starts out on, so you can play the instrument from your computer's keyboard straight away; turn it off when you would rather your keystrokes didn't generate musical accompaniment.

Below the Synthesizer page's panel sits the modulator strip, and below that the macro row. Both are described later in this chapter.

### Working the Controls

Drag a knob up or to the right to turn it up, and down or to the left to turn it down, or use your mouse wheel or trackpad over it. Hold SHIFT while you drag or scroll for finer adjustment, at a third of the usual speed. Either way, the display at the top of the window shows the knob's value as it changes. Double-click a knob to return it to its default value. Right-click a control to open its modulation popup, which is described under Modulation Routing below; on an iPad, press and hold it instead. You can choose either gesture on any platform with the Open Modulation Panels setting, described under Settings below.

Drag the three dots at the bottom-right corner of the window to resize it. The window keeps its proportions as it grows and shrinks to avoid squashing the contents. On an iPad there is no corner to drag. Sandbender grows as large as the space it is given allows without changing its proportions, centered in that space. If the space is more than twice as wide as it is tall, Sandbender fills the width instead, and you scroll up and down by dragging one finger on any part of the panel that isn't a control.

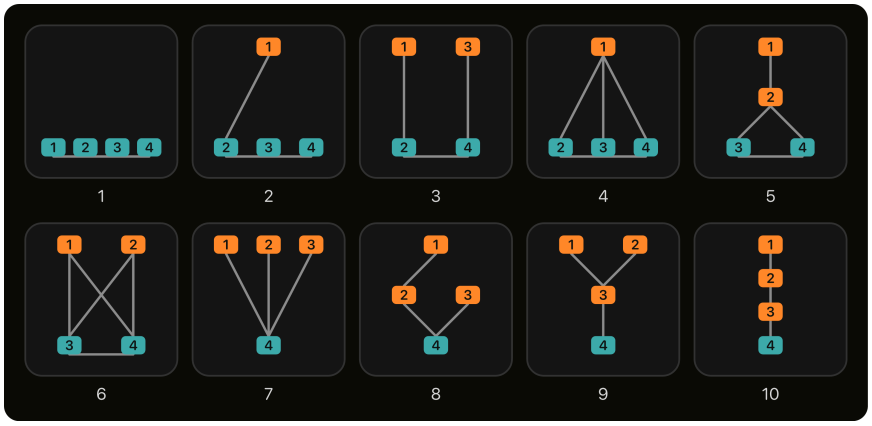
## Operators

The operators section presents four identical columns, one per oscillator, with the section's shared controls in two columns of their own to the right.



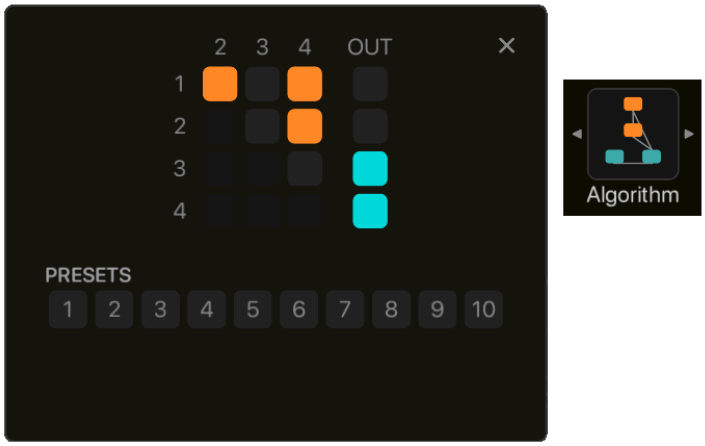

### Algorithm

The ALGORITHM control, in the column just to the right of the operators between DETUNE and FOLD, chooses how the four operators are wired to one another. The small arrows beside it step through the pre-wired factory algorithms, wrapping around from the last to the first, and the preset display shows the number of each one as you land on it. There are ten: the first nine are shamelessly lifted from the RYK module, and the tenth adds the full serial stack with all four operators in a chain.




In addition to the small display in the control itself, the operator section draws the current algorithm for you. An orange dotted line runs from each operator that modulates another to the operator(s) it feeds, a gold dotted line joins the operators you can hear, and a cyan dot under an operator's LEVEL knob marks it as audible. The lines brighten when you point at a FREQ knob of an operator they connect to. The ring around each LEVEL knob tells the same story in color: gold for an operator you hear, orange for one that modulates another, red for one that does both, and grey for one that, in the current algorithm, does neither.

Clicking the ALGORITHM control itself opens a larger view in which you can rewire the operators. The column of numbers 1–4 represents the operators as sources; the row of numbers 2–4 and the word OUT represent destinations. The squares represent all possible connections between sources and destinations, observing these restrictions: An operator may feed any higher-numbered operator; it may not feed any lower-numbered operator, and it may not feed itself. There are six such connections available, and any combination of them is legal. Also, any operator can be made audible by clicking its switch below the OUT heading. Clicking a grey square turns on the connection; clicking it again turns it off. Turn every operator's OUT switch off and the operators fall silent, though the sub oscillator and the noise still sound.




The larger view stays up to date while it is open: step the arrows, load a preset or let your host automate the algorithm, and the view follows. Any changes you make to the algorithm are saved with the preset like any other setting.

**Per-Operator Controls**

Each of the four columns carries the same three controls. From top to bottom:

LEVEL sets the operator's output, from silence to full. It's important to understand that this knob does two jobs at once: an operator's LEVEL scales both how loudly it is heard *and* how hard it modulates. There is no separate modulation depth. This is how the DX-style instruments worked and we have kept it, because decoupling the two levels would turn a playable instrument into a spreadsheet. If operator 1 is wired to operator 2, turning up LEVEL 1 makes operator 2's tone brighter and more complex, and if operator 1 is also audible it gets louder at the same time. The default values put operator 1 at full level and operator 4 at 80%, since operator 4 is the carrier in most of the factory algorithms. The middle two default to silence, waiting to be dialed in.

FREQ sets the operator's frequency, and what that means depends on the MODE dots beneath it:

- RATIO: the operator tracks the keyboard, and FREQ selects a whole-number harmonic relationship to the note you played. This is the usual FM arrangement and the one that gives stable, pitched results.
- FREE: the operator still tracks the keyboard, but FREQ is a continuous multiplier from 0.5 to 18 rather than a list of harmonics. Ratios between the whole numbers are where inharmonic, bell-like and metallic sounds live.
- FIXED: the operator stops tracking the keyboard altogether and sits at an absolute frequency, anywhere from 0.5Hz to 18kHz. A fixed operator modulating a tracked one gives you the formant-like sounds that move in timbre but not in pitch as you play up the keyboard.

WARP is the control that makes this section more than a four-operator FM chip. It is a single bipolar knob that morphs the operator's waveform continuously: centered it is a pure sine, turned anti-clockwise it moves toward a triangle, and turned clockwise it moves toward a pulse. Nothing switches; the shape changes continuously. Because a modulator's waveform determines the sidebands it produces, WARP on a modulator is a timbre control of enormous range, and WARP on an audible operator is simply a tone control. Try a two-operator algorithm, put both operators at the same ratio, and sweep WARP on the modulator while listening to the carrier.

There is one more per-operator control that has no knob on the panel: PITCH, an offset of plus or minus 24 semitones. It exists as a modulation destination, so that an envelope or a sequencer can move a single operator's tuning without touching the others. To reach it, route a source to the operator's FREQ knob: modulation aimed at FREQ always moves PITCH, whichever of the three modes the operator is in, so a vibrato means the same thing on a FIXED operator as on a RATIO one.

**Section Controls**

The two columns to the right of the operators hold the controls that act on all four. The left column carries REGEN, DETUNE, the ALGORITHM control described above, and FOLD; the right column carries XMOD, SPREAD and, next to FOLD, dots for choosing the fold type.

XMOD is a global scaling on every modulation connection in the algorithm at once, from nothing to full. Because it multiplies what the LEVEL knobs are already doing, it is the fastest way to open and close the whole timbre of a patch, and it is the first thing worth assigning to an envelope or a macro. At zero, no frequency modulation happens and every operator is a bare warped oscillator no matter how the algorithm is wired.

REGEN is feedback: operator 4 modulating itself. It is the one connection the algorithm view cannot draw, because every other connection runs from a lower-numbered operator to a higher one, and so it gets a knob of its own. Operator 4 gets it because

operator 4 is the one operator that can never modulate anything else — there is nothing above it to feed — which is exactly why the four-operator Yamaha chips of the 1980s put their feedback loop in the same place. At zero there is no feedback. Turned clockwise, the operator's sine leans over into something like a sawtooth, getting brighter and buzzier as it goes; at the top of the travel it collapses into noise. The musical part of the range is roughly the bottom third, and the rest is there for when you want the noise. The feedback is scaled by operator 4's own LEVEL, as all of its modulation is, so turning that LEVEL down calms the feedback as well as the volume.

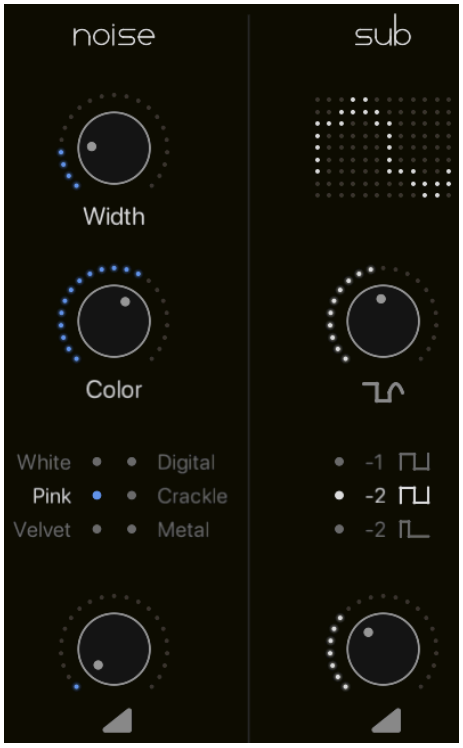
DETUNE spreads the pitch of the four operators against one another. At zero they are in exact ratio; turning it up moves them apart by increasing amounts, which thickens a patch and, past the middle of its travel, turns a pitched sound into a cluster. It acts on the operators, not on the voices, so it thickens a single note as readily as a chord.

SPREAD fans the audible operators out across the stereo field. At zero they all sit in the center. Turning it up moves them apart, evenly and in operator order, until at full the lowest-numbered audible operator is hard left, the highest-numbered is hard right, and any others are spaced between them. Only the operators the algorithm makes audible take part; any heard only as modulation have nowhere to be panned to. If only one operator is audible, SPREAD has nothing to spread, and does nothing. The sub oscillator and the noise source are not panned by it. SPREAD works best on algorithms with two or more carriers, and with a little DETUNE it turns a single note into something with real width.

FOLD and the fold-type dots beside it control a wavefolder on the audible sum of the operators. FOLD is the drive into it: at zero the folder does nothing, and turning it up pushes the signal through progressively more folds, adding harmonics that have nothing to do with the algorithm. The three modes differ in how they fold:

- SYMMETRIC is the classic reflecting folder, mirroring the signal about the upper and lower limits. It is even-handed and adds odd harmonics.
- ASYMMETRIC folds the positive half of the signal at one threshold and the negative half at half that, which is a rougher, brassier sound.
- SOFT CLIP does not fold at all. It rounds the signal off as it approaches the limit, which is saturation rather than folding. It's much gentler, and the one to choose when you want FOLD to add weight rather than transform.

## Sub




The sub oscillator is a divider hung off operator 1, in the way that the sub oscillators of the classic monosynths were hung off their single oscillator: it is clocked by operator 1's cycle and is therefore always exactly in tune with it, whatever operator 1 is doing.

TYPE chooses the division: one octave down as a square, two octaves down as a square, or two octaves down as a pulse. The pulse setting is thinner and sits differently in a mix; the square settings are the ones that add oomph.

The WARP knob (the one with the squiggly symbol) crossfades the sub's shape from square toward a triangle of matching duty cycle. A square sub is a blunt instrument that will dominate anything above it; turning this knob up softens it into something that supports a patch instead of burying it.

LEVEL is the sub's contribution to the bus. It starts at zero.

Because the sub follows operator 1 rather than the note, putting operator 1 into FIXED mode leaves the sub sitting at a fixed division of that frequency too, turning it from a sub oscillator into a second drone.

## Noise

The noise source is the other channel on the bus, and it is rather more than the usual hiss generator.

TYPE chooses the character: WHITE and PINK are the familiar two. VELVET is sparse impulse noise, which sounds like noise in aggregate but is mostly silence, and is interesting fodder for a resonant filter. DIGITAL is quantized and gritty. CRACKLE is intermittent, in the manner of a worn record. METAL is tuned and ringing, the one to use when you want noise that has a pitch.

COLOR tilts the noise darker or brighter around its center position.

WIDTH spreads the noise in stereo: at zero both channels carry the same signal, and at full they are two uncorrelated streams. The sub is always mono, and the operators are mono unless SPREAD pans them, so a little noise WIDTH under an otherwise centered patch is a cheap and effective way to give it some air.

LEVEL is the noise's contribution to the bus, and it too starts at zero.

## Filter




The filter is a bank of eleven topologies, one per voice, sitting between the operator sum and the output stage. TYPE selects among them, or OFF, which bypasses the filter entirely.

- SVF LP, SVF HP and SVF BP are a clean state-variable filter in its three responses. This is the neutral choice, handy when you want a filter to gently shape the sound without imposing its own personality.
- LADDER LP, LADDER BP and LADDER HP model the four-pole transistor ladder, which loses bass as resonance rises and is the sound of a great many records.
- DIODE LP is the diode ladder, the harder-edged relative of the transistor ladder.
- SK LP and SK HP are Sallen-Key, a two-pole design that goes from gentle to screaming across its resonance range.
- K35 LP and K35 HP are the filter from a certain Japanese monosynth of the late seventies, which distorts as you drive it in a way that probably dismayed its designers but is its reason for inclusion here.

FREQ is the cutoff, from 20Hz to 20kHz, and it defaults to the top of its range, which means the filter is effectively open until you close it. Q is resonance, and how far you can take it before the filter self-oscillates depends entirely on which topology is selected. DRIVE pushes the signal into the filter; turning this up on the ladder and K35 types brings out their character.

## Output Stage

Underneath the filter is the output stage, which is where the amplitude envelope meets the audio and where a voice acquires its shape in time. The dots choose between three behaviors.

- VCA is a plain amplifier. The amplitude envelope controls loudness and nothing else. A square envelope gives you a square note. This is the predictable choice and it is the default.
- LPG GATE and LPG BOTH are a model of the Buchla 292 low-pass gate. Inside a low-pass gate is a vactrol — a lamp glued to a photoresistor — and the envelope does not control the audio directly. It controls the lamp, and the photoresistor follows the lamp with a lag that no envelope generator can produce and that you cannot dial out, because it is a physical property of the cell. Struck with a short envelope, a low-pass gate rings down like a plucked or struck object rather than fading like a knob being turned.

In GATE mode the vactrol acts as an amplifier only. In BOTH mode it acts as an amplifier and a low-pass together, which is what the circuit really does: as the note dies away it also gets duller, because the photoresistor's resistance rises on both paths at once. BOTH is the characteristic West Coast sound, and it is the one that makes short envelopes sound like struck metal and wood.

Two controls sit with the output stage. VELOCITY sets how much MIDI velocity affects the stage, from none to full. It is not specific to the low-pass gate: it scales the VCA's gain and the gate's control voltage alike, so a hard strike lands the same way whichever stage you have selected. Because the low-pass gate's brightness follows its level, velocity on an LPG mode opens the tone as well as raising the volume, which is exactly what you want and is very hard to achieve with an envelope alone.

RESPONSE is the speed of the vactrol's cell, from fast to slow around a center position, and it does nothing at all with the plain VCA. Turned down, the gate is quick and percussive; turned up, it smears the strike and lets notes bloom and overlap.

One more knob sits in this part of the panel, between DRIVE and RESPONSE, marked with a small symbol rather than a name: LEVEL, which your host lists as Pre-FX Level. It sets how loud the signal is going into the effects rack, from silence at fully anti-clockwise to unity at fully clockwise, which is where it starts. It does a similar job to the MASTER knob in the macro row; the two are multiplied together, and both act before the effects rather than after them. However, they are not the same. MASTER is the patch's overall volume, sits at 80% by default, and cannot be modulated. LEVEL can: route an LFO to it for tremolo across the whole instrument, or an envelope to duck the dry sound into a delay. Because it comes before the effects, turning it down also drives the rack less hard, which matters to Kombinat and Evil Otto, while the delay and reverb tails ring on after LEVEL has pulled the dry sound away.

Unlike most destinations, LEVEL is applied to all eight voices together. If a modulation source runs separately in each voice, e.g., the envelope generators, the most recently started voice that is still sounding, release and all, is the one that drives it.

## The Modulator Strip

Under the panel runs a horizontally scrolling rail carrying the instrument's modulation sources. From left to right: a playable keyboard, the MIDI input, two envelopes, two function generators, two LFOs and a step sequencer. Drag the rail sideways, by any part of it that isn't a control, to reach the ones that are off-screen, or use the buttons in the macro row, described below, to jump straight to one. The same rail appears under the Effects page; the effects are legitimate modulation destinations.




Every module has one or more output rows, and each row carries a small circular token on its right, a name, and an activity LED that shows what that output is doing. The token is how you patch: drag it onto any control in the instrument and you have made a connection. Modules collapse when you double-click at the top of their pane if you would rather have the panel space, and which ones are open is saved with the preset.

Some of the modules' knobs are marked with a small picture rather than a name. Where they are, this manual calls them by the name displayed in your host's parameter list.

### Modulation Routing

Anything in the strip can be routed to very nearly anything in the instrument, including the controls of other modulators.

To make a connection, drag a source's token from the modulator strip onto any control. The connection is made, and the control's modulation popup opens so that you can set its depth straight away. To see a control's connections later, right-click it, or press and hold it, depending on the Open Modulation Panels setting. The popup lists everything currently routed to it, one row per source.

The popup works from either end. Opened on a control, it shows that control's incoming connections. This is the fastest way to find out why something is moving when you did not ask it to. Opened on a source's token, it shows everything that source is wired to, one row per destination.

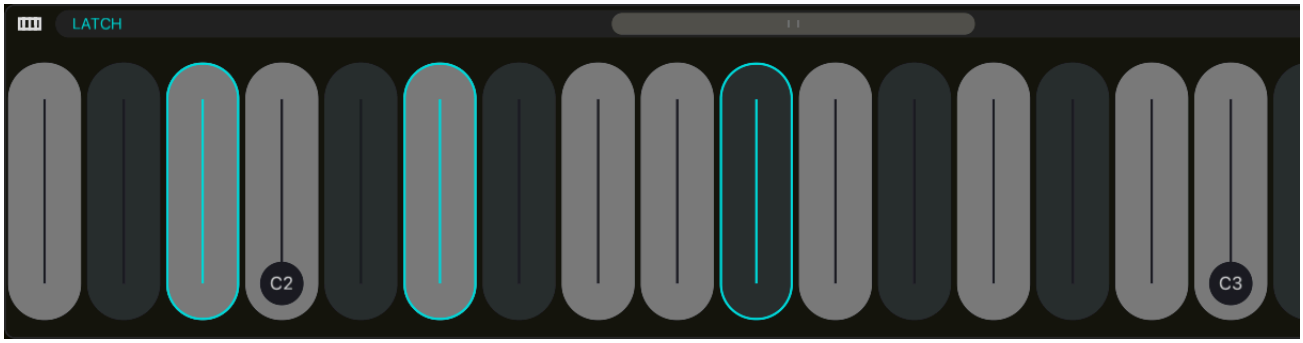
Opened on a control, the popup's top line shows the control's name and its value, and the value is a text box: type a new one and press RETURN, or click elsewhere, to set the control. Below that, each row carries the other end's name, an X to remove it, and a bipolar depth slider with a detent at its center. If set to the right of center the source pushes the destination up; to the left it pulls it down. A source routed at zero depth is connected but doing nothing, which can be useful when you're refining a patch.




We'll describe each of the modulation sources next.

**Keyboard**

The first module on the rail is a keyboard you can play with the mouse or, on a touchscreen, with up to ten fingers at once. It is there in every version of Sandbender, not just the standalone, and it is how you audition a patch when there's no MIDI controller within reach.




All 128 MIDI notes are available, with each C labeled. The keyboard opens at C3. C3 is MIDI note 48, so middle C is C4. The strip along the top of the module moves you around the keyboard: drag it sideways to scroll, or up and down to make the keys bigger or smaller. The arrows beside it do the same in steps, an octave or a zoom size at a time.

The keys respond to more than just being pressed. Where you strike a key sets the note's velocity: at the top of a key it is full, and it falls away toward the bottom. Once a note is down, keep your finger or the mouse moving. Sliding up and down sends MIDI controller 74 (MPE's usual expression controller) measured from wherever you first touched. Sliding sideways bends the note's pitch, by one semitone for every key's width you travel. With MPE turned on in Settings, every note bends on its own, up to an octave either way. With MPE off, the bend is the instrument's pitch wheel, limited to two semitones, and every note sounding follows it. Either way, a note you let go of mid-bend stays at the bent pitch while it dies away rather than sliding back.

LATCH holds notes after you let go of them, so you can build a chord and take your hands away to work on the patch. Press a held key again to release it, or turn LATCH off to release them all. Loading a preset releases them as well.

The keyboard has no modulation outputs of its own. Its notes and control messages go into Sandbender by the same path MIDI notes do, so they show up on the MIDI input module's rows exactly as if you had played them on a controller. The traffic runs the other way too: notes arriving from your host or a MIDI controller light up the keys they play, on every channel, so you can watch what an MPE controller or a sequencer track is sending. In the standalone version, notes you play on the computer keyboard are treated the same way.

**MIDI Input**

The second module carries the modulation sources that arrive as MIDI. Each is a row of its own, with its own token:

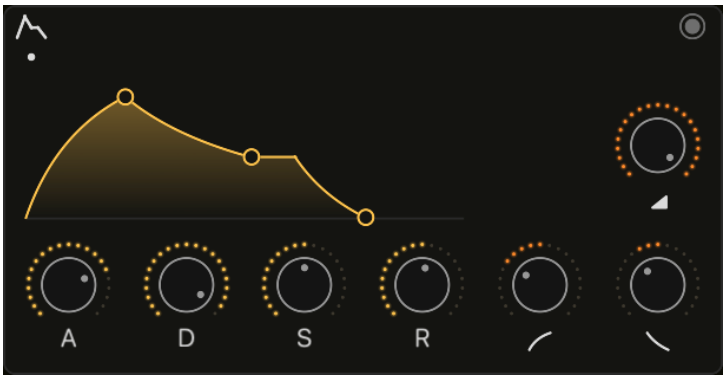
- NOTE is key tracking: the note's pitch, as a modulation value.
- VELOCITY is how hard the note was struck.
- PRESSURE is aftertouch, i.e., how hard you lean on held keys.
- PITCH BEND is the pitch wheel, as a modulation source in its own right, separate from the pitch bend that the instrument applies to the note anyway.
- MODULATION is the mod wheel.
- The five remaining rows are the assignable continuous controllers, each labeled with the controller number it listens to. You choose those numbers on the Settings page.

With MPE turned on, PRESSURE, PITCH BEND and MODULATION belong to each note individually rather than to the instrument as a whole, and so does any controller row set to 74 or 1. A note's own controller 74 and controller 1 are measured from the first value that note receives, so they swing on either side of zero from wherever your finger landed, which is how MPE slide is meant to behave. That is also how you put MPE slide to work: leave the first controller row on 74, its default, and route that.

**Envelope Generators**

Envelope generators shape continuous sounds into discrete notes whose shape changes over time. Sandbender's EGs are of the standard ADSR variety with a couple of extra features.

ENV 1 is the amplitude envelope. It is wired directly to the output stage and needs no routing to do its job, though nothing stops you routing it elsewhere as well. ENV 2 reaches the instrument only through connections you make.




ATTACK runs from 0.1 milliseconds to 10 seconds, DECAY and RELEASE from 1 millisecond to 20 seconds, and SUSTAIN is the level the envelope holds while a key is down. These are the usual features of an ADSR and shouldn't be any surprise to you.

To add a bit more flexibility, we have two more knobs for altering the shape of the envelope. ATK CURVE bends the attack segment, from exponential through linear to logarithmic across its travel, with linear at the center. FALL CURVE does the same for the decay and the release *together*, on the reasonable grounds that they are the same gesture happening at different times and a patch almost never wants them shaped differently. Both are bipolar and both default to center.

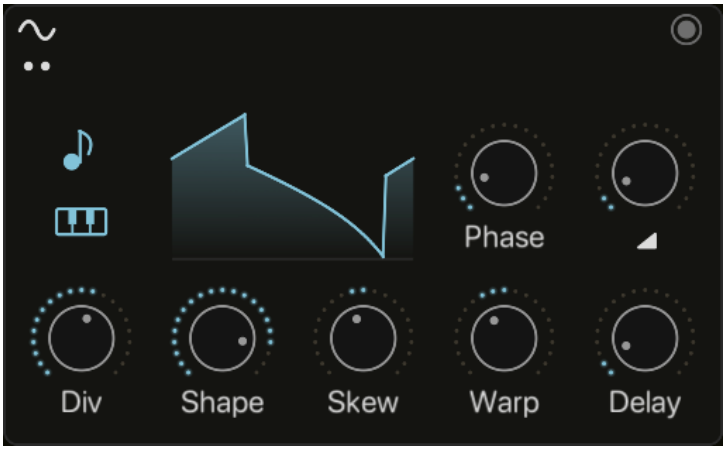
One refinement is worth knowing about if you like your attacks sharp. Bent hard enough, an attack curve finishes almost vertically, and a short attack that ends vertically is a click. To compensate, the attack is never allowed to climb faster than twice its average rate. With ATK CURVE at center this never comes into play, and it only starts to matter once the knob is more than about a third of the way from center toward either end. Turned anti-clockwise, where the attack leaps away from silence and flattens into its peak, the limit softens the start of the climb and the attack still finishes on time. Turned clockwise, where the attack eases out of silence and arrives steeply, it softens the arrival instead, and at the extreme the attack can finish up to about a sixth later than the ATTACK time says.

What the limit can't do is make a very short attack gentle. It is measured against the ATTACK time, so an attack of a fraction of a millisecond is still a jump of a few samples, and it can click at any curve setting. If a patch clicks in a manner you don't like at the start of every note, give the envelopes a millisecond or two of attack.

LEVEL scales the envelope's whole output. On ENV 1 this changes the overall volume of the voice; on ENV 2 it's the depth control that affects every connection you've made to the envelope generator's output.

**LFOs**

The two low-frequency oscillators are the same flexible design we use in Quanta, and their output varies smoothly through a very large number of shapes rather than switching between a handful.




RATE runs from 0.01Hz (one cycle every hundred seconds) to 50Hz, which is well into the audio range and useful for grit rather than movement. Turn SYNC on and RATE is replaced by DIVISION, a metrical value from 2/1 (one cycle every two bars) down to 1/32, with D and T marking dotted and triplet values.

Four controls shape the wave, and their effects are related enough that watching the display while you turn them will teach you more than this paragraph can. SHAPE morphs the basic waveform smoothly through four familiar ones: a double sine, a sine, a triangle and a square. SKEW adjusts the wave's horizontal symmetry, which on a triangle takes you from a falling ramp through a

symmetrical triangle to a rising ramp, and on a square varies the duty cycle. WARP applies curvature, bending a triangle into a spike or rounding a sine into a pulse. PHASE sets where in the cycle the wave begins.

SYNC and RETRIGGER are the two small switches on the module, marked with a musical note and a keyboard respectively, lit when they're on, grey when they're off. Your host lists their two settings as FREE and SYNC, and FREE and NOTE.

RETRIGGER decides whether the LFO belongs to the note or to the host's timeline. Turned off, the LFO is one oscillator running continuously for the whole instrument, and every voice reads the same value at the same moment. With SYNC on as well, that oscillator locks to your host's beat whenever the transport is running, so the LFO's peaks land in the same place in the bar every time through; with the transport stopped it carries on at the tempo-derived rate. If RETRIGGER is on, each voice gets its own copy of the LFO, restarted at the PHASE setting whenever a note begins. The two cannot both be true: a beat-locked phase belongs to the timeline and a retriggered phase belongs to the note, so with RETRIGGER on, SYNC sets only the rate. In MONO voice mode, a note played legato over another doesn't restart a retriggered LFO, for the same reason it doesn't restart the envelopes.

DELAY fades the LFO in over a set time after each note starts, from nothing up to five seconds, for the classic way of getting vibrato that arrives after the note rather than with it. The fade belongs to each note whether RETRIGGER is on or off. Finally, LEVEL scales the output.

**Function Generators**

The two function generators are the most modular thing in the instrument. Each is, in effect, one channel of a Maths[6] Eurorack module, and each is simultaneously three things: a function generator, an LFO and a slew limiter. It turns out that all three are the same circuit with a different target, a clever insight for designing an analog module and just as useful in a software counterpart.




RISE and FALL set the two segment times, each from half a millisecond to thirty seconds. Once triggered, the channel rises over the RISE time, and when it arrives at the top it falls over the FALL time. That is an attack-release envelope.

Turn CYCLE on and the generator does not stop: it rises, falls, rises again, and you have an LFO whose two halves are independently timed, which is something no shape knob will give you.

Route a modulation source into the channel's IN jack and the behavior changes entirely. The channel now chases that input rather than a fixed value, at the rate RISE and FALL allow. What you have is a slew limiter. This is how you smooth a stepped source, tame a fast envelope, or turn a square LFO into a trapezoid. You do not select this mode; it happens because something is patched into the input, exactly as it would on a hardware module.

RESPONSE is the curve of both segments, bipolar around a linear center. There is one detail about it worth reading twice: the fall retraces the rise. With RESPONSE turned down, the channel moves quickly at low values and slowly at high ones, in *both* directions, so the rise is fast off the bottom and flattens into its arrival at the top, and the fall is slow off the peak and accelerates on its way down. That is not the ADSR convention, where a decay is fast off the peak, and next to the envelope display it looks wrong. This is a deliberate, accurate model of the hardware's behavior in both directions.

MODE, the switch marked with a keyboard, decides whether playing a note triggers the channel. Lit, it is TRIG, and each note sets the channel rising from wherever it happens to be, then falling, then stopping unless CYCLE is on. Unlit, it is FREE, and notes are ignored altogether: the channel moves only if CYCLE is on or something is patched into it. FREE does not by itself make the

[6] Product name is a trademark of Make Noise Co.

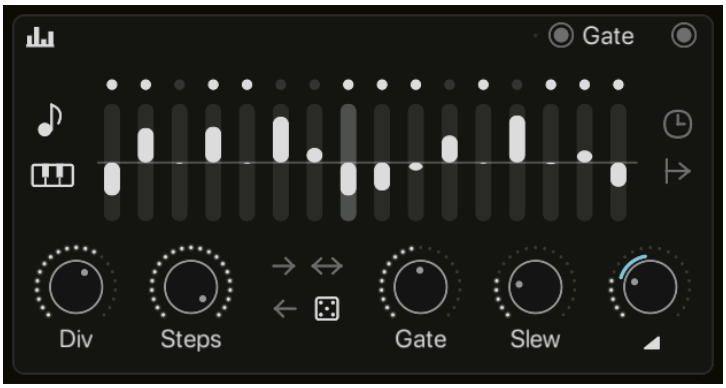
channel run; CYCLE does that. A channel cycling in FREE is one oscillator for the whole instrument, just like an LFO with RETRIGGER off: every voice reads the same value at the same moment. There are two exceptions. Something patched into the IN jack makes the channel a slew limiter rather than an oscillator. A single shared oscillator can't follow an input that is different for every note, so a channel with an envelope, velocity, key tracking, the sequencer, a retriggered LFO, a function generator on TRIG or MPE expression patched into any of its inputs runs separately in each voice instead. The TRIG jack, described below, works in either mode, since a patch cord is not the keyboard.

Each channel publishes two outputs. The first is its value, as you'd expect. The second is its gate, labeled EOR for End Of Rise, which is high for as long as the channel is falling. Patch that gate into the other channel's TRIG jack and you have two function generators chasing each other, a good start for interesting patches.

Three of a channel's six inputs are jacks with no associated knob: IN, BOTH and TRIG. They exist only to be the destination of a connection. The other three inputs need no jack, because a control already stands where the jack would be: route something at the RISE or FALL knobs to get rise and fall CV, and route something at the CYCLE switch to get a cycle input, which the engine combines with the switch exactly as the hardware combines the button with its jack.

**Step Sequencer**

Sandbender's sequencer is a classic analog step sequencer rather than a groove box: sixteen stages, each with a value and a gate, plus a clock, a length, a direction and a reset. There is no swing, no ratcheting and no per-step probability, and their absence is a deliberate choice of usability over complexity.




The two lanes are independent, giving you much of the functionality of two sequencers in one. The value lane and the gate lane advance together, but a step with its gate off is a rest for the gate lane *only*, while the value lane still lands on that step and still outputs its value. Hence you can do things like gating an envelope on four steps out of sixteen while the value lane sweeps continuously underneath.

Each step's value runs from −100% to +100%, centered at zero, so a sequence can push a destination both ways from wherever its knob sits. Each step's gate is simply on or off. Click in the value lane to set a step to the height you clicked at, or drag across the lane to draw a shape through several steps. Since the exact middle is hard to hit by hand, double-click a step to set it back to zero. Click a gate to turn it on or off, and drag across the gates to set a whole run of them the way the first one went.

RATE sets the internal clock, from 0.05Hz (twenty seconds a step) to 50Hz, and SYNC and DIVISION replace it with a metrical value in the way the LFOs do. LENGTH sets how many of the sixteen steps are used, from one to sixteen. DIRECTION runs the sequence forwards, backwards, alternating, or randomly. GATE sets how much of each step the gate lane is high for, and so how long the notes are. LEVEL scales the value lane's output.

MODE, the switch marked with a keyboard, decides what a new note does to the sequence. Lit, it is TRIG, and every note starts the pattern again from its first step, or the last step when DIRECTION runs backwards, or a random one when it runs randomly. Unlit, it is FREE, and notes leave the pattern running wherever it is. Each voice has its own copy of the sequencer, so in FREE the voices of a chord drift out of step with one another, which is the point. Either way the sequence loops for as long as the note lasts; there is no one-shot mode.

SLEW is worth a word. Every other modulation source in the instrument is continuous, and a step sequencer is the first thing here that jumps. A full-scale jump aimed at an operator level or a filter cutoff will click, so SLEW rounds the corner off the value lane (and only the value lane, since a gate must stay a gate). It defaults to a deliberately non-zero two milliseconds, which is short enough to still sound stepped and long enough to stop the clicking. Turn it up, as far as a full second, for a sequence that glides between steps.

Two jacks sit on the module. CLOCK takes over the timing: route anything at it — a function generator's gate, an LFO, another part of the patch — and the internal clock stops, with every step coming from a rising edge instead. As with the function generators'

slew mode, this is not a setting you select. RESET returns the sequence to its first step. The RATE knob needs no jack, because routing something at the knob is the same thing.

## The Macro Row

The strip along the bottom of the Synthesizer page holds the four macros, a row of buttons for the modulator strip, and the controls that belong to the instrument as a whole. From left to right: MACRO 1 to 4, the modulator buttons, PORTAMENTO, the voice mode dots, the RANGE dots, TUNE, MASTER, and the output meters.




**Macros**

Each of the four macro knobs is a source with no behavior of its own: it is a knob you can route to as many destinations as you like, each with an independent depth, so that one gesture moves a whole patch.

They are routed exactly like any other source: drag the macro's token onto a control, and set the depth in the popup that opens. Macros are ordinary parameters, so your host can automate them, and they are the sensible thing to put on a hardware controller. The PERFORM surface, described later in this chapter, can also play them by hand.

Name a macro by clicking its label, typing, and pressing RETURN. The names are saved in the preset.

**The Other Controls**

Between MACRO 4 and PORTAMENTO sits a row of six small buttons, one for each kind of module on the modulator strip: the keyboard, the MIDI input, the envelopes, the function generators, the LFOs and the sequencer. Click one and the strip opens every module of that kind and slides along to the first of them; click it again once you're there and it collapses them. Like the macros' tokens, they appear only on the Synthesizer and Effects pages, where the strip does.

The voice mode dots choose how Sandbender allocates its voices; click a dot or the name beside it. POLY is the usual arrangement, with up to eight notes sounding at once. MONO and RETRIG both play one note at a time, and if you hold several keys, the most recent one wins; let go of it and the instrument falls back to the one you were holding before. The difference between these two modes is in the envelopes. In MONO, a note played while another is still held slides into it without starting the envelopes again, which is legato in the monosynth sense: the envelopes start only when you play from silence. In RETRIG, every note starts the envelopes afresh.

PORTAMENTO is the time the pitch takes to glide from one note to the next, from off at fully anti-clockwise up to two seconds. It is a time rather than a rate, so a glide across an octave takes as long as one across a semitone. It only applies in MONO and RETRIG, since in POLY each note gets its own voice and has nowhere to glide from. Most of the knob's travel is given to the short times, which are the ones you'll use most.

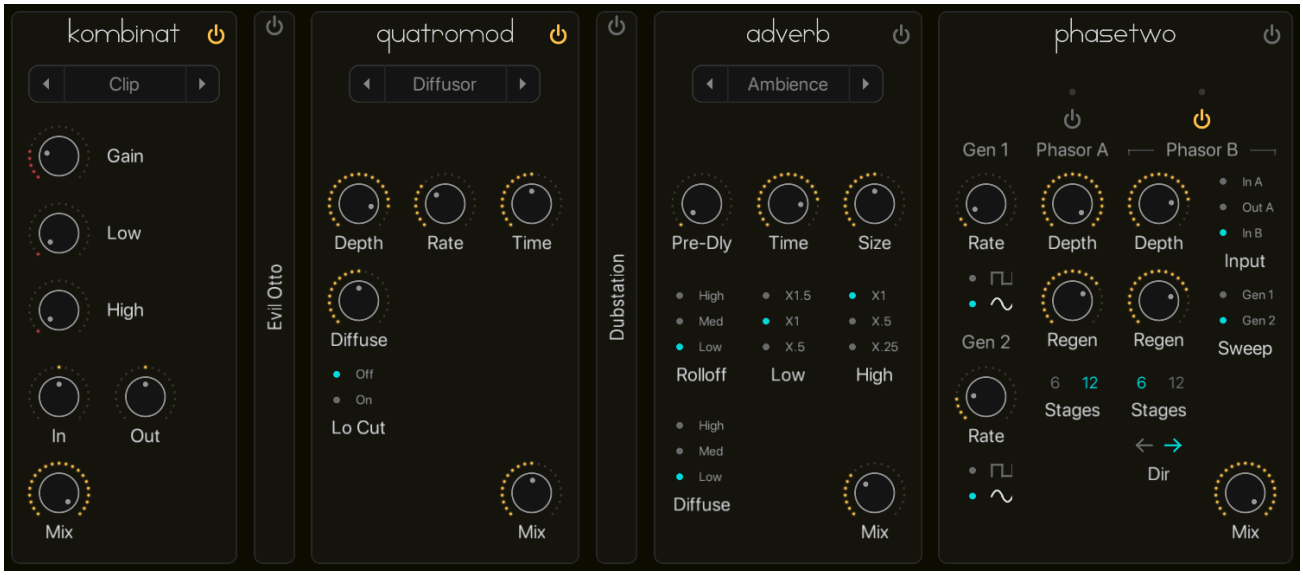
The RANGE dots transpose the whole instrument in octaves, marked in organ footages from 32′ down to 2′, with 8′ as the center position.

TUNE is fine tuning for the whole instrument, plus or minus 100 cents.

MASTER is the patch's overall volume, starting at 80%. It is applied before the effects, together with the LEVEL knob in the output stage; that section explains how the two differ. The meters to its right show the peak level of the left and right channels leaving Sandbender, after the effects.

## Effects

The Effects page holds a rack of six effects in series after the synthesizer. Each is built with the same DSP code as the Audio Damage plugin it is named after, running inside Sandbender.




Drag an effect by its name to move it within the chain. The order is saved with the preset, which is important because, for instance, a distortion before a reverb and a distortion after one produce different sonic results. Double-click the name to collapse an effect into a narrow strip, hiding its controls. Double-click the strip to return it to its usual width.

With every effect open, the rack is wider than the window, so it scrolls sideways like the modulator strip. Use the scroll bar beneath it, drag any part of the rack that isn't a control or an effect's name, or use your mouse wheel or trackpad. While you are dragging an effect to a new position, pushing it toward either edge of the window scrolls the rack for you.

Two controls are present in every effect. The power button to the right of the name takes the effect out of the chain altogether. The MIX knob sets the balance between the signal that went in and the signal that came out, from fully dry to fully wet.

The modulator strip runs under this page as well, so every effect knob is a modulation destination like any other. Note, however, that some modulators run within Sandbender's voices while the effects run after the voices. If a modulator restarts with every note (like the envelope generators), the modulator in the most recently started voice that is still sounding will control the effect's parameter. When no note is sounding at all, an envelope routed to an effect does nothing, while the LFOs, the function generators and the sequencer carry on from instrument-wide copies of their own.

We'll provide brief descriptions of all of the effects and their controls here. This should be enough information to get you tweaking the effects without further ado. If you're interested in more in-depth descriptions of these effects, you can find the manuals for the original plugin versions online here: https://github.com/Audio-Damage-Inc/manuals

### Evil Otto

A multiband upward-and-downward compressor of the sort that makes everything sound bigger and slightly dangerous. It splits the signal into three bands and squeezes each toward a target from both directions at once.

IN and OUT trim the levels on either side of the processing, plus or minus 12dB. DEPTH sets how hard the effect works. TIME sets the attack and release of all three bands together, from fast and pumping to slow and gluing. DOWN and UP set how much downward and upward compression is applied respectively, and turning UP down is the difference between a compressor and the characteristic sound this effect is known for. LOW, MID and HIGH are band trims after the processing, plus or minus 12dB each.

### Dubstation

Our modeled bucket-brigade delay, still wildly popular even though it's old enough to vote.

TIME L and TIME R set the two delay times independently, from 4 to 1000 milliseconds. Turn SYNC on and they are replaced by DIVISION L and DIVISION R, metrical values on the same table the LFOs use. Independent left and right times are always on, so setting them to different values is how you get the classic dub scatter.

REGEN is feedback, from −30dB to unity. At the top of its range the delay does not quite run away, letting you ride it to just this side of disaster.

LO CUT and HI CUT filter the repeats, from 100Hz to 1.5kHz and from 4kHz to 8kHz respectively. Narrowing them is what makes each repeat sound older than the last. SATURATION dirties the feedback path, and a little of it with the cuts brought in gets the bucket-brigade character.

RATE and DEPTH control an LFO modulating the delay time, from a slow drift to an obvious warble. WIDTH sets the stereo spread of the repeats. LOOP freezes the delay line, holding whatever is in it and ignoring new input. This is a performance control more than a mix tool.

**ADverb**

One of our algorithmic reverbs, with all three of the algorithms we added over time.

ALGORITHM chooses between PLATE, AMBIENCE and HALL, which are three genuinely different reverbs rather than three sizes of one. PRE-DELAY holds the reverb off for up to 200 milliseconds, which is how you keep a big reverb from swallowing the attack of what fed it. TIME is the decay, from 0.1 to 9.9 seconds. SIZE scales the space itself, from 1 to 99.

DIFFUSION, with choices of LOW, MED or HIGH, sets how quickly the early reflections smear into a wash; low settings keep individual echoes audible and suit percussive material. ROLLOFF, also with a three-way switch, sets how much high end the tail loses as it decays.

LOW and HIGH are multipliers on the decay time at the two ends of the spectrum. Low contour runs ×0.5, ×1 and ×1.5, and high contour ×0.25, ×0.5 and ×1, so you can have bass that rings on after the highs have faded, or the reverse.

**Quatromod**

Four modulation effects in one panel, chosen with a mode popup menu. Each mode keeps its own settings, so changing modes and returning finds everything where you left it, and the panel's dials adjust themselves to whichever mode is selected.

FLANGER gives you two controls for the sweep LFO: RATE from 0.02 to 10Hz and DEPTH from zero to max. MANUAL sets the static delay offset on either side of zero. OFFSET adds stereo spread between the left and right sides. REGEN adjusts the feedback amount, with INVERT flipping the polarity of the feedback and turning the sweep hollow instead of resonant.

CHORUS gives you DEPTH and RATE modulation controls, TIME for the base delay, and REGEN for feedback. Its MODE switch chooses between FLUID and ENSEMBLE, providing one voice moving smoothly or a thicker multi-voice arrangement.

DIFFUSOR is somewhere between a chorus and a small reverb but does not sound quite like either. DEPTH and RATE modulate a diffusion network rather than a delay line, TIME sets its length, DIFFUSION sets how dense it is, and LO CUT keeps the bass from going crazy. (Some people consider the bass going crazy to be a good thing about this effect, so try it both ways.)

FREQ SHIFT is a frequency shifter, not a pitch shifter: DEPTH and RATE modulate the shift, AMOUNT sets it from −100% to +100%, REGEN feeds the output back in, and STEREO runs the two channels in the same direction or in opposite ones. Small amounts create a phaser that never repeats; large amounts are inharmonic and strange.

**Kombinat**

Fourteen distortion algorithms, from mild to hostile: THRU, FUZZ, SATURATE, DISTORT, CLIP, WARP, DEGRADE, RINGMOD, X NOISE, OCTAVER, WAVY, RECTIFY, TUBE CLIP and NERD RAGE.

ALGORITHM chooses among them, and the three knobs beneath it do whatever that algorithm says they do, and the panel relabels them as you switch. FUZZ and SATURATE use only AMOUNT; DEGRADE gives you RATE, BITS and ERROR; RINGMOD gives you AMOUNT, FREQ and SHAPE; OCTAVER gives you DRY, 1/2 and 2/1 to balance the dry signal against octaves below and above. THRU passes the signal untouched, which is useful when you want the input and output trims and nothing else.

IN and OUT are trims of plus or minus 24dB. On most of these algorithms the input trim is the real control: how hard you hit it decides what it does.

**Phase Two**

This plugin started as an accurate emulation of the Mu-Tron Bi-Phase and grew some features of its own as we refined it over the years. It has two independent phasors and the two sweep generators that drive them. The interesting sounds come from the two phasors moving independently in their own ways.

GEN 1 and GEN 2 are the sweep generators. Each has a RATE, from 1 to 10 in the plugin's own units rather than hertz, and a SHAPE switch between SQUARE and SINE. A square sweep steps the phaser between two positions rather than gliding, which is a sound in itself.

Phasor A has ACTIVE to switch it in, STAGES at 6 or 12, DEPTH, and REGEN for feedback. It is always driven by Generator 1.

Phasor B has the same ACTIVE, STAGES, DEPTH and REGEN, and three more controls. SWEEP chooses whether Generator 1 or Generator 2 drives it. Put both phasors on Generator 1 and they move together; put B on Generator 2 and they drift against each other, which is the most interesting aspect of the original hardware. DIRECTION runs B's sweep NORMAL or REVERSE, so the two

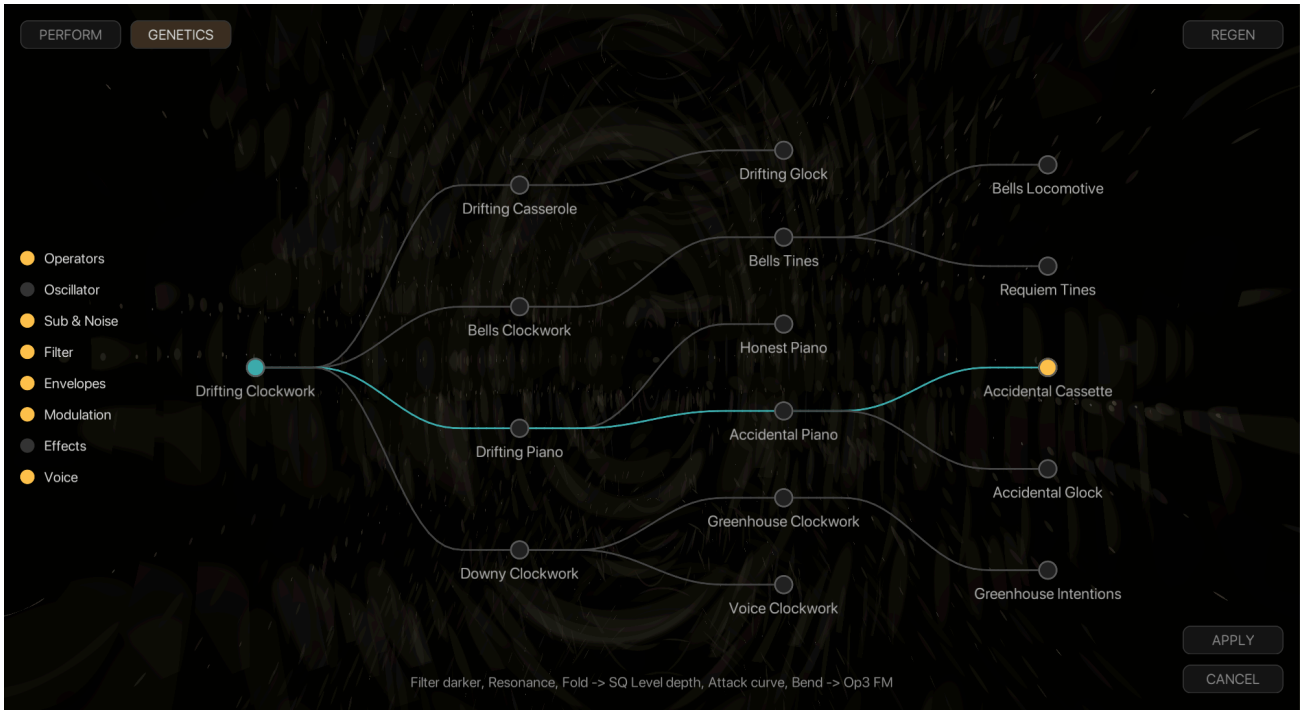
phasors can move in opposite directions from the same generator. INPUT decides what B actually processes, and with it where the two phasors end up in the stereo field. Phasor A always processes the left input and always feeds the left output; B feeds the right output, and INPUT chooses where it gets its signal:

- IN B, the default, feeds B from the right input. The two phasors work side by side on the two channels, which is the true stereo arrangement and the one to use on a stereo source.
- IN A feeds B from the left input, the same signal phasor A is getting. The two phasors run in parallel on one source, and the stereo image comes entirely from the difference between them. Put B on GEN 2 or set its DIRECTION to REVERSE and the sweep chases itself from side to side.
- OUT A feeds B with phasor A's output, so the right side hears the two phasors in series and the left side hears A on its own.

Under IN A and OUT A, the right input is ignored.

## Genetics

Sandbender covers a wide sonic territory, and the next interesting sound you find is often a long way from where you start. The Genetics page is a way of exploring that territory without having to know in advance where you're going. It takes the sound you have, breeds a family of variations on it, and lets you audition them one at a time until something catches your ear.




Click the GENETICS icon at the top left of the window, then the GENETICS button, to begin. Sandbender starts with the current settings, unsaved changes and all, as the ancestor of the family, and grows a tree of descendants from it: three or four children, each with one or two children of its own, producing a final generation of four to eight great-grandchildren scattered among them. The ancestor sits at the left, with a blue node, and each generation lies further to the right.

Every descendant is a copy of its parent with some changes. The further a sound is from the ancestor, the further it has evolved. A child might differ from the ancestor in ten or so ways, and a great-grandchild in forty or more, since it carries its parents' and grandparents' changes as well as its own. The changes aren't completely random knob-twiddling. Each sibling in a family is given a different theme, such as its operators, its modulation or its effects, and makes changes with that theme, so siblings tend to go off in different directions. Under the tree, a line of text lists the main changes in whichever descendant you've selected.

Click any node in the tree to hear the child's sound. The change happens immediately, without interrupting notes you are holding, so the easiest way to explore is to hold a chord, or latch one on the keyboard, and click around. (This assumes, of course, that the sounds have some sustain to them. You'll need to hit new notes if you're exploring percussive or plucked sounds.) The node you've selected turns gold, and the path back to the ancestor lights up in blue.

Double-click a sound to make it the ancestor of a new tree. This is how you go further: when you find something promising, double-click it and its descendants will explore the territory around it. REGEN, at the top right, goes back to the original ancestor — the sound you had when you clicked GENETICS — and grows a fresh tree from it.

When you've found a sound you want to preserve, click APPLY. The tree display closes and the sound stays, renamed after the node you chose. It is not saved yet: it is a new patch with no preset file behind it, so save it from the preset browser if you want to keep it. Click CANCEL instead and the tree display closes and the original sound comes back.

If you close the overlay by clicking GENETICS again, or simply switch to another page, Sandbender keeps whichever sound you last clicked but doesn't change the name of the preset. If you load a preset while the overlay is open, that preset becomes the new ancestor, and CANCEL will take you back to it rather than to what you had before.

The column of round switches on the left chooses what Genetics is allowed to change: OPERATORS, OSCILLATOR (the algorithm and the section controls), SUB & NOISE, FILTER, ENVELOPES, MODULATION (the modulators and their routings), EFFECTS, and VOICE (voice mode, portamento and the output stage). All of them start switched on; turn off the ones you want left alone. The change takes effect the next time a tree grows, with REGEN or a double-click.

A few things are never changed, whatever the switches say: MASTER, the output stage LEVEL, TUNE, the level of envelope 1, the macros, the order of the effects, the PERFORM loop, and everything on the Settings page. A PERFORM loop that is playing carries on moving its nodes while you audition, and a control held by a node stays where the node has it, whatever a descendant does to its knob. Genetics will rewire the algorithm, add, remove and rescale modulation routings, and switch effects in and out, but it works within limits. Every descendant has something audible coming out of it, stays within about 30% of the ancestor's loudness, keeps its fundamental pitch and the speed of its attack and release, and has its brightness, feedback, folding, resonance, reverb and delay kept on a leash. It is designed to surprise you, not to deafen you.

Some changes only show up when you play. A new routing from velocity or key tracking won't be heard until the next note, and one from the mod wheel, aftertouch or a macro won't be heard until you move it.

## Perform

Turning knobs one at a time is a fine way to build a sound but a clumsy way to play one. The PERFORM surface gives you two handles, called nodes, to slide around the window, each of which moves two controls at once, and records what you do with them as a loop that runs in time with your song.




Click the PERFORM button at the top left of the Genetics page to open the surface, and click it again to put it away. The surface fills the page, with the eye candy still swirling behind it and the macro row still in place beneath it. Sandbender doesn't remember that PERFORM was open: close and reopen the window and you are back on the ordinary Genetics page.

### The Nodes

The two nodes are the numbered rings on the page. Each one is a point with two axes: where it sits from left to right is its X axis, and where it sits from bottom to top is its Y axis. Along the bottom of the surface are four menus, labeled 1X and 1Y on the left and 2X and 2Y on the right, and each one connects an axis to a control. The choices are the four macros, XMOD, REGEN, FOLD, DETUNE,

SPREAD, each operator's LEVEL and WARP, the sub's LEVEL, the noise's LEVEL and COLOR, the filter's FREQ, Q and DRIVE, and the MIX of each of the six effects. An axis can also be aimed at NONE, should one axis be enough for your artistic expression.

Touch the page and the nearer node jumps to your finger; drag, and it follows. With two fingers you can move both nodes at once, which is where the surface comes into its own on an iPad or a touchscreen. With a mouse or a trackpad you move one node at a time: click, and the node nearer the pointer comes to it.

A node doesn't turn the knobs it is aimed at; it modulates them. Each knob stays where it is, and its modulation arc runs out to wherever the node has put the control, from the knob's minimum at the left or bottom edge of the field to its maximum at the right or top. Let go, and the control stays there. In fact it stays there until you move the node again, aim the axis at something else, or load a preset, and turning the knob in the meantime changes nothing you can hear, because the node is holding the control where it is. (Any other modulation routed to the control still works on top of the node.) To hand a control back to its knob, set its axis's menu to NONE or to another control.

An axis aimed at a macro works the same way. The macro's own knob stays put, and everything the macro is routed to follows the node, exactly as if you had turned it.

Until a node takes hold of a control, the axis follows it instead: aim an axis at a control and the node moves to show where that control's knob is. Touch the page, though, and the node comes to your finger, taking the control with it.

A node's ring is gold while you are holding it or the loop is playing it, and grey at rest. While it is gold, a note beside it lists its two axes, what each one is aimed at, and that control's value as the node has it.

**The Loop**

Along the bottom of the surface runs the transport for the loop recorder. A word at its left says what the loop is doing, a progress line across its top is marked off in bars, and four small squares at its right, one for each axis, fill in when the loop holds something for that axis. In the middle are the main button, a STOP button and the loop's length, with the four axis menus on either side of them.

The loop is between one and eight bars long, and starts out at two. The arrows on either side of the length change it, but only while the loop is stopped, and changing the length empties the loop, since what was in it was played against the old length. The bars are your host's bars, in its time signature. The standalone version has no song to follow, so there the loop runs at 120 BPM in 4/4.

The main button's symbol shows what the loop is doing, or, when it is stopped, what a press will start:

- On an empty loop, it starts recording. Recording runs for one pass of the loop and then switches to playback by itself; press the button during the pass to switch to playback sooner.
- On a stopped loop with something in it, it starts playback.
- During playback, it switches to overdub, and pressing it again switches back.

Presses land on the beat. A press waits for the next bar line — your host's bar line while your host is playing, the loop's own otherwise — and pressing again before then takes it back. The exception is a stopped loop under a stopped host, which has no bar line to wait for and starts at once. STOP doesn't wait: it stops the loop immediately. To empty the loop, press and hold STOP until the red arc around it winds all the way around.

While your host plays, the loop follows its song position, so it stays in step when you jump around in the song. When the host stops, the loop doesn't; it carries on at the last tempo it saw. It also keeps running when you close Sandbender's window, so press STOP when you've heard enough.

Recording replaces the existing loop. Holding a node records both of its axes, and as the recording pass goes by, the axes of a node you aren't holding are wiped from the loop, so a single pass holds exactly what you did during it and nothing else. With two fingers you can record both nodes in one pass. With a mouse, record one node and then overdub the other: in OVERDUB, the node you are holding is recorded over for as long as you hold it, and the other carries on playing.

The loop records where you put each axis, not what the controls were doing, so aim an axis somewhere else and the same gesture plays on the new control. Playback moves the nodes, and the nodes modulate their controls exactly as your hand did. Take hold of a node during playback and you have it for as long as you hold it; in PLAY that isn't recorded, and when you let go the loop has it back. When the loop stops, each node stays where the loop left it, and so do its controls.

**Saving**

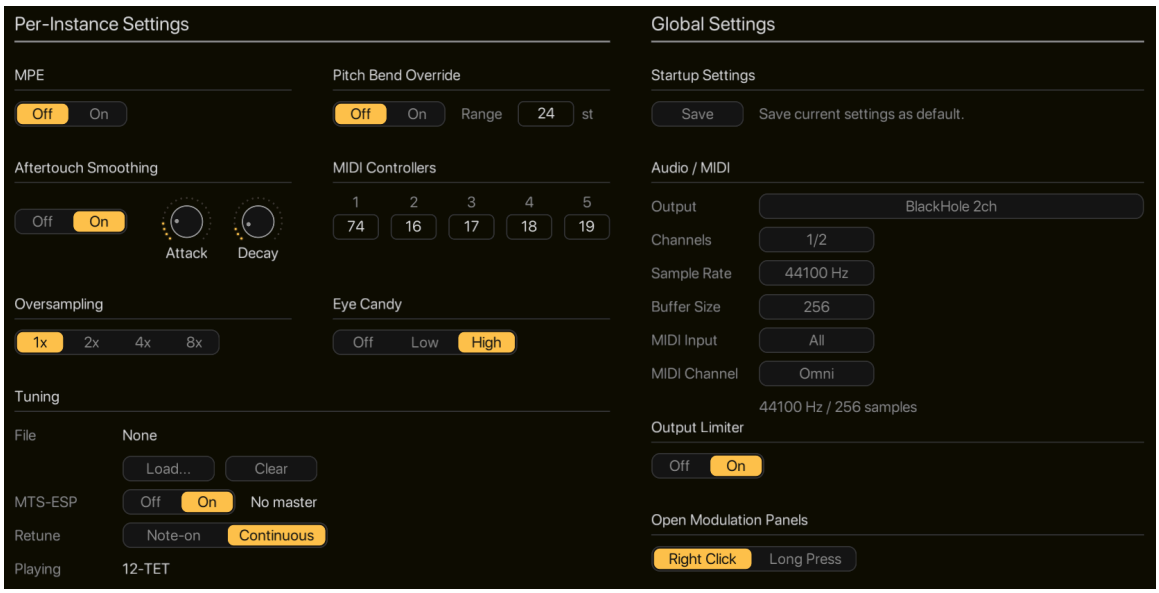
A preset saves what the four axes are aimed at, and the loop if there is one. When the preset loads, the loop is stopped and ready to play, and every control a node was holding is let go, back to its knob, until you move a node or play the loop; where you last left the nodes isn't saved. Neither the loop nor the axes' targets are saved with your host's project, so save a preset to keep them.

Loading a preset replaces both, and that includes auditioning presets in the browser. CANCEL in the browser brings back your sound but not your loop, so save a loop you care about before you go browsing.

PERFORM is invisible to your host. The nodes modulate; they don't change any parameter's value. Nothing they do, whether by hand or from the loop, is recorded as automation, and neither the loop nor the axes are parameters you can automate.

## Settings

The Settings page holds preferences that belong to the instrument rather than to a preset. Nothing here is saved with a patch, so loading a preset never changes them.




The page is divided in two. The left-hand block, PER-INSTANCE SETTINGS, belongs to this particular copy of Sandbender: change something there and only this instance is affected, and your host saves it with the rest of the session. The right-hand block, GLOBAL SETTINGS, holds the controls that reach beyond one instance. We'll take the left-hand block first.

### Tuning

As you'd expect, Sandbender defaults to 12-note equal-tempered intonation, since this is what the vast majority of contemporary musicians and their instruments use. If you're thinking, "Why would I use anything else?" you can probably skip the rest of this section. On the other hand, if you regularly work with other scales and tunings, you'll be pleased to see that Sandbender has several relevant features.

#### MTS-ESP

Sandbender utilizes the MTS-ESP session-wide tuning system created by ODDSound, Ltd. https://oddsound.com/. To the right of the MTS-ESP label are two buttons, OFF and ON. If you click the ON button, once an MTS-ESP master plugin is present in your DAW, Sandbender will connect to it. The name of the current scale is displayed to the right of the word PLAYING.

MTS-ESP is ON by default. With it on, a word beside the buttons tells you what Sandbender has found: CONNECTED when there is a master in the session, NO MASTER when there isn't. Sandbender also understands tunings sent to it as MIDI Tuning Standard system-exclusive messages; if one arrives while there is no master, the word reads SYSEX and Sandbender plays in that tuning. On iOS, where MTS-ESP isn't available, the row doesn't appear at all.

There are two buttons to the right of the RETUNE label: NOTE-ON and CONTINUOUS. If you choose CONTINUOUS, Sandbender automatically updates its tuning whenever any changes are made in the master. NOTE-ON does what it says: changes to the scale made by the master only take effect when a new note starts. CONTINUOUS is the default.

#### TUN Files

Sandbender's pitch reference, intonation, and/or keyboard mapping can be completely changed with TUN files. TUN files are a standard format for providing tuning information. If you're familiar with Scala files but not TUN files, the Scala program can export TUN files. One TUN file contains the information that Scala places in separate SCL and KBM files. A web search will reveal plenty of information about creating TUN files. The presence of MTS-ESP takes priority over the TUN file.

To load a TUN file, click the LOAD… button. Sandbender will prompt you with a standard file dialog box; choose your file and Sandbender will adopt the tuning and mapping found in that file. Once a file is loaded, Sandbender keeps a reference to that file

within its settings. If you move the TUN file, you'll have to reload it into Sandbender. To revert to the usual 12-note equal-tempered tuning, click the CLEAR button.

**MIDI Controllers**
Sandbender receives up to five assignable continuous controllers, which appear as the five controller rows of the MIDI input module. Each of the fields numbered 1 to 5 takes the CC number you want the corresponding row to listen to; type a number from 0 to 127 and press RETURN. The defaults are the ones a controller is most likely to already be sending: 74, which the MPE specification uses for slide, and 16 through 19, the general-purpose block.

These are settings rather than parameters. A preset does not get to remap your hardware.

**Oversampling**
The engine runs at 1x, 2x, 4x or 8x the host's sample rate, defaulting to 1x. Higher settings reduce the aliasing that a wavefolder and a hard-driven operator bank inevitably produce, and cost proportionally more CPU. 1x is the lightest setting, which is why it is the default. If a bright patch sounds gritty in a way you do not want, raising this control to 2x or 4x is usually enough. Conversely, if Sandbender is heavier on the CPU than you would like, check that this hasn't been left on a high setting.

**Eye Candy**
The animated backdrop on the Genetics and About pages, and behind the PERFORM surface, at OFF, LOW or HIGH. It is a storm of particles that reacts to what you play: notes stir it up and add to it, controllers ruffle it, and it brightens with the instrument's output. LOW draws about half as many particles as HIGH, and OFF hides it altogether. It's purely cosmetic, and worth turning down on an older machine.

**MPE**
The MPE switch affects how Sandbender handles MIDI notes and controller messages. If you have an MPE controller such as a Roger Linn Design LinnStrument or Madrona Labs Soundplane, turn this switch on. Sandbender will then process pressure, pitch bend, and position information independently for each note, providing the expressivity that you expect from your controller. If you don't have an MPE controller, leave this switch turned off and ignore it. MPE off is the conventional arrangement: every note arrives on the same channel, and pitch bend and pressure apply to everything sounding.

As you're probably already aware, you also need a host program that can correctly transmit multi-channel MIDI information from your MPE controller to Sandbender. Some hosts have some sort of switch that needs to be turned on to enable MPE with virtual instruments.

Sandbender uses MPE's lower zone: channel 1 carries messages for the whole instrument, and channels 2 through 16 carry one note each. With MPE on, polyphonic aftertouch goes to the note it names; with MPE off, Sandbender treats it as ordinary aftertouch for the whole instrument. Unlike everything else on this page, the MPE switch is a parameter your host can automate.

**Pitch Bend Override**
The MPE specification says that a note's own pitch bend covers 48 semitones either way, and that is what Sandbender assumes. Not every controller agrees: many can be set to a narrower range, which gives you finer control of the pitch. If yours is set to something other than 48, turn PITCH BEND OVERRIDE on and type your controller's range, from 0 to 48 semitones, in the RANGE field beside it, so that a given movement of your finger bends the note by the amount your controller expects.

The override applies only to the per-note bend that MPE sends. The ordinary pitch wheel always bends Sandbender by two semitones either way, whether MPE is on or off.

**Aftertouch Smoothing**
Pressure from a controller can arrive in coarse steps or jitter as you hold a key, and routed to something sensitive, like the filter's cutoff or an operator's level, that roughness is audible. Turn AFTERTOUCH SMOOTHING on, and the two dials beside it slow the pressure down: ATTACK sets how long the value takes to follow pressure that is rising, and DECAY how long it takes to follow pressure that is falling. Each runs from 0 to 1000 milliseconds, with most of the dial's travel given to the shorter times. A short ATTACK and a longer DECAY is a good place to start: the sound responds as soon as you press, and relaxes gracefully when you let up.

The smoothing applies to all pressure, whether it arrives as ordinary aftertouch for the whole instrument or as a note's own pressure under MPE.

**Startup Settings**
The first item in the right-hand block is a SAVE button labeled "Save current settings as default." Every new instance of Sandbender starts with the settings it saves: MPE, the pitch bend override and its range, aftertouch smoothing and its times, the five controller numbers, oversampling, eye candy, MTS-ESP, RETUNE, the TUN file, and the output limiter. Changing a setting on

the left of the page affects only the instance you are working in; SAVE is how you make it the starting point for the instances you create from now on. Instances that already exist keep the settings they have. The two settings SAVE leaves alone are Open Modulation Panels and Tooltips (see below), which are saved the moment you change them.

**Audio / MIDI**
This section appears only in the standalone versions of Sandbender. In a plugin, your host takes care of all of this.

DRIVER chooses the audio system, and appears only on computers that offer more than one. OUTPUT chooses the audio device; on iOS, where the system decides where audio goes, it simply tells you. On a computer, CHANNELS then chooses which pair of the device's outputs Sandbender plays through, which is handy with an interface that has more than two. Choosing a new OUTPUT sets it back to 1/2. SAMPLE RATE and BUFFER SIZE do what they say, and a smaller buffer gives you a quicker response to what you play in exchange for more work for your computer. MIDI INPUT chooses the MIDI device Sandbender listens to, or ALL of them, and MIDI CHANNEL chooses whether it listens to all sixteen channels (OMNI) or just one. MPE needs all sixteen, so MIDI CHANNEL is ignored while MPE is on. Beneath the menus, a line of text confirms the sample rate and buffer size actually in use, or tells you what went wrong.

These choices take effect immediately and are remembered the next time you run Sandbender; the SAVE button above doesn't need to be involved.

**Output Limiter**
The last thing Sandbender does to its output, after the effects, is pass it through a limiter. It is fast-acting and has no lookahead, so it adds no latency. Anything that peaks above −1dB is pulled down to it, and a gentle saturation behind the limiter catches the very front of a transient that gets through before the limiter can react, so nothing leaves the plugin louder than −0.1dB.

It is ON by default, because an FM instrument with a feedback knob, a wavefolder and a rack of effects can produce some astonishing peaks, and we'd rather you heard them at a sensible level. If you are deliberately driving something after Sandbender, or would rather manage the level yourself, turn it OFF. Although it sits in the right-hand block, it is saved with each instance, like the settings on the left.

**Open Modulation Panels**
This chooses how you open a control's modulation popup: RIGHT CLICK, or LONG PRESS, which means pressing a control and holding it still for a moment. Only the one you choose works, so a right-click does nothing on LONG PRESS, and holding a control does nothing on RIGHT CLICK. (On a Mac, a CTRL-click counts as a right-click.) It starts out on RIGHT CLICK on a computer and on LONG PRESS on an iPad. On an iPad, which has no right button, that is where it should stay. LONG PRESS is also the one to choose if you play Sandbender on a touchscreen.

This is a preference of yours rather than of any one instance. It applies to every copy of Sandbender on the machine, is saved the moment you change it, and isn't part of what SAVE stores.

**Tooltips**
Rest the mouse pointer on a control for a moment and a box beside it gives the control's name and a sentence or two about what it does. Move on to the next control and its tooltip comes up sooner, so you can read your way along a row of knobs without waiting at each. This switch turns them OFF or ON; they start out ON. Like Open Modulation Panels, it applies to every copy of Sandbender on the machine, is saved the moment you change it, and isn't part of what SAVE stores. An iPad has no pointer to rest anywhere, so there are no tooltips there, and no switch.

## Automation

Every control on the Synthesizer and Effects pages is a parameter your host can automate. The Settings page is the exception: of everything there, only the MPE switch can be automated. The rest are deliberately excluded. The MIDI controller numbers are excluded for the reason given above, OVERSAMPLING because changing it changes the delay Sandbender reports to your host, and the others because they describe your setup rather than your music. The function generators' and sequencer's input jacks, which exist only as the far ends of modulation routings, aren't automatable either. The modulation routings themselves are saved with the preset but are not automatable.

The four macros are ordinary parameters, and are ideal automation targets. A macro moving eight destinations at once is one lane in your host instead of eight. PERFORM, on the other hand, is invisible to your host: its nodes modulate their controls rather than turning them, and neither the nodes nor the loop can be automated.

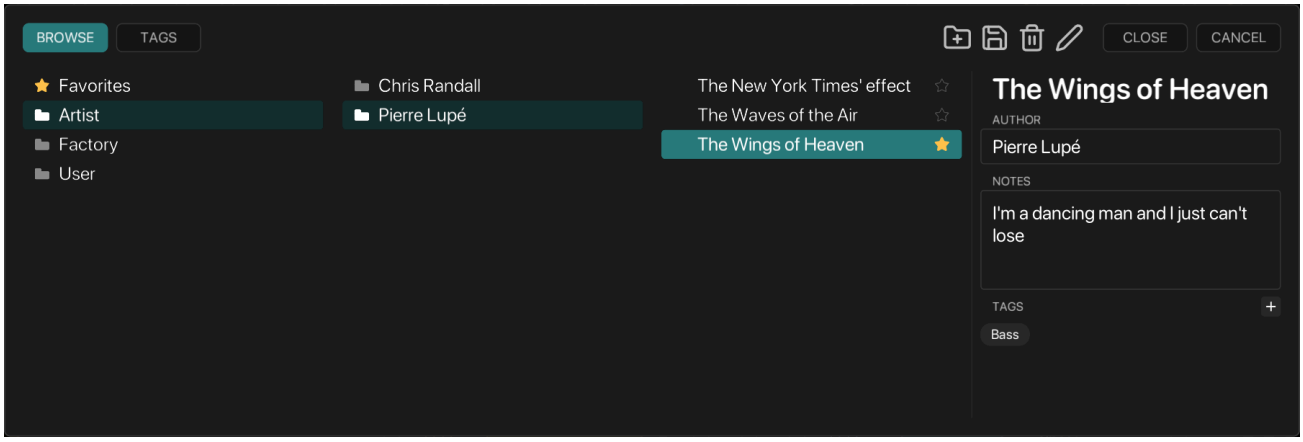
Sandbender has a great many parameters, and scrolling through them all in a host's list is no fun. In the CLAP version, Sandbender offers your host eight pages of remote controls — banks of eight parameters each, which hosts such as Bitwig Studio map straight onto a hardware controller's knobs:

| Page | Controls |
|---|---|
| Main | MACRO 1–4, filter FREQ, filter Q, XMOD, MASTER |
| Freq & Level | FREQ and LEVEL for operators 1–4 |
| Shape | WARP for operators 1–4, FOLD, REGEN, DETUNE, fold type |
| Sub, Noise, Output | sub LEVEL and WARP, noise LEVEL, COLOR and WIDTH, filter DRIVE, RESPONSE, VELOCITY |
| Amp Envelope | ENV 1 ATTACK, DECAY, SUSTAIN, RELEASE, ATK CURVE, FALL CURVE, PORTAMENTO, output stage LEVEL |
| Envelope 2 | ENV 2 ATTACK, DECAY, SUSTAIN, RELEASE, ATK CURVE, FALL CURVE, LEVEL |
| LFOs | RATE, SHAPE, SKEW and WARP for both LFOs |
| FX Mix | MIX for all six effects, Dubstation REGEN, ADverb TIME |

The FREQ controls on the Freq & Level page set the operators' RATIO values, whatever mode the operators are in.

## Presets

Sandbender keeps its presets in a browser that floats over the middle of the window. Click the preset name in the display at the top of the window to open it.




### Browsing

Two tabs sit along the top of the browser: BROWSE, which shows your presets as folders and files, and TAGS, which shows them by description. We'll take BROWSE first.

The browser displays presets and folders in scrollable lists, arranged in columns. The leftmost list shows the top of Sandbender's preset collection: FACTORY, which holds the presets that come with Sandbender, and USER, which is where your own work goes. FACTORY is divided by kind of sound: Ambient, Basics, Bass, Dubstep, FX, and so on. Once you have marked at least one preset as a favorite, as described below, a FAVORITES folder appears above them.

Click a folder and its contents appear in the next column. Folders can contain folders of their own, as deep as you like: each time you open a folder, the columns slide one place to the left to make room for its contents, and once there are more columns than fit, a scroll bar appears so you can slide back. To the right of the columns is a detail pane, which shows the preset you have selected: its name, a banner for its folder if the folder has one, and its AUTHOR, NOTES and TAGS. You can type into the AUTHOR and NOTES fields, and click the + beside TAGS to add or remove tags, and your changes are written straight into the preset file.

Click a preset's name to hear it. The browser stays open, so you can click your way down a list auditioning as you go. Double-click a name to load the preset and close the browser in one gesture.

**Loading a preset replaces whatever you were playing**, including any changes you haven't saved, so if you have created something you want to use again, save it as a new preset before loading another.

Once you have clicked on any item in the browser, you can navigate with your computer's keyboard. The left and right arrow keys move between columns, and the up and down arrow keys move within a list, auditioning each preset as you land on it. RETURN loads the selected preset and closes the browser, just as a double-click does.

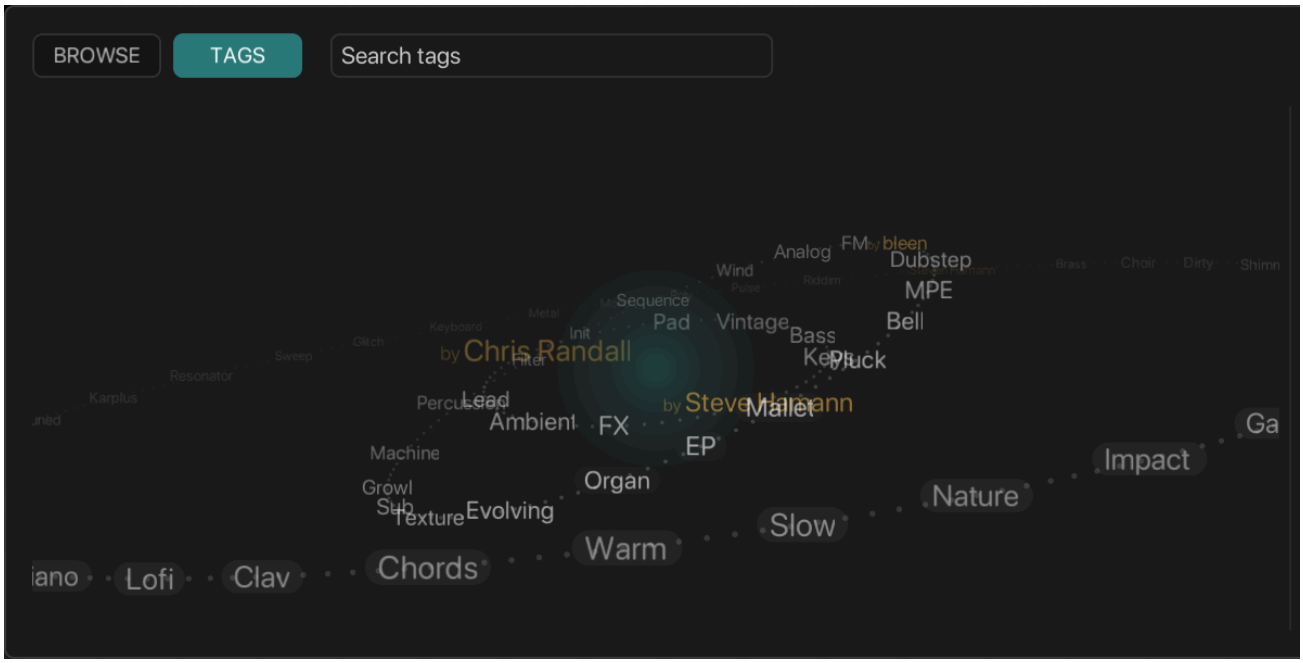
There are two ways to close the browser, and the difference matters. CLOSE, or a click anywhere outside the browser, closes it and keeps whatever you last auditioned. CANCEL, or the ESC key, closes it and puts back the sound you had when you opened it, unsaved changes and all. That makes it safe to go browsing in the middle of working on a sound: if nothing you hear is better than what you had, CANCEL takes you straight back.

The stars to the right of the preset names mark presets as favorites. Click an empty star to mark a preset, and click a filled star to remove the mark. The star in the preset display at the top of the window does the same for the current preset. Once you've marked at least one favorite, the FAVORITES folder appears in the leftmost column; clicking it shows you everything you have marked, gathered from wherever it happens to live, and clicking a name loads it as usual. FAVORITES is a list rather than a real folder, so you can't save, move or drop anything into it.

The arrows on either side of the preset display step through every preset in the order the browser lists them, folder by folder, without your having to open the browser at all. They skip FAVORITES, and past the last preset they wrap around to the first, and vice versa.

## Tags

Folders are good for knowing where a preset is; tags are good for knowing what it is. Each preset can carry any number of tags — "bass", "bell", "pad", "evolving", and so on — and the TAGS tab lets you find presets by them.



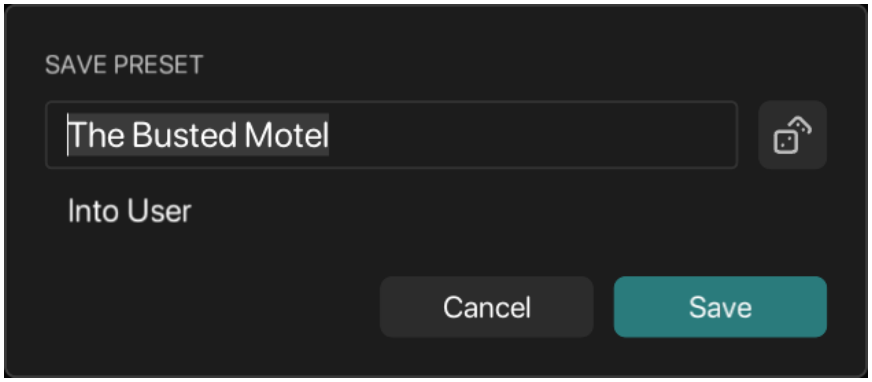

The TAGS tab shows every tag in the collection as a cloud, along with the name of every preset author, marked with a "by" and drawn in a color of their own. The words circle the middle of the pane on tilted rings, seen in perspective, so each one swings around in front and passes behind, larger and brighter the nearer it is, with the most-used words largest of all. Drag anywhere in the cloud to turn the whole thing; flick it and it coasts to a stop. Click a tag or an author and it moves to the center, and the list beside the cloud shows every preset that carries that tag or was made by that author. The tags and authors that share presets with it rearrange themselves around it, closest first. Click a second one and the list narrows to the presets that match both, so you can ask for one author's bass sounds, say; click anything in the center to let it go again. The search field on the tab row narrows the cloud to tags containing whatever you type. Clicking a preset in the list loads it, exactly as in BROWSE.

To change a preset's tags, select it in BROWSE and click the + beside TAGS in the detail pane. The cloud reappears with the preset's tags highlighted; click a tag to add it or remove it, or click the + at the end of the cloud to type a new one.

## Saving and Managing Presets

Four buttons at the right-hand end of the tab row manage the collection: NEW FOLDER, SAVE, DELETE and RENAME. Each acts on whatever you have selected: the selected preset or folder, or, if nothing in a column is selected, the folder that column is showing. Each asks you to confirm, or to type a name, before it does anything.

To save a preset, select the folder you want it in, or a preset in that folder, and click SAVE. The SAVE PRESET box tells you which folder the preset is going into and offers you a name. If you had selected a preset, the name is that preset's: leave it as it is to save over that preset, or type a new name to save alongside it.




The dice at the right-hand end of the name field invent a name for you. These are not random words: the generator listens to the patch — how bright it is, how gritty, how long it rings, how high or low it sits, how much space the effects give it, how much it moves — and makes up something that suits it. Roll again if you don't like the first one. The dice do their best to avoid names that are already in the folder.

DELETE removes the selected preset or folder from your storage device after you confirm, and it is gone for good. RENAME renames it. You can also delete with your computer's DELETE key.

You can move presets and folders by dragging and dropping them. Drag a preset onto a folder in any column to move it into that folder, or drop it onto the open space of a column to move it into the folder that column is showing.

## Copying and Pasting Presets

Preset files are plain text, so you can share them in forums, email them to your friends, and move them between computers, whatever the operating system. The browser makes this easy: select a preset and use your system's copy command (CTRL-C on Windows and Linux, CMD-C on macOS) to put the whole preset on the clipboard as text. Select a folder and use the paste command to turn text on the clipboard back into a preset in that folder. Sandbender checks that what you are pasting really is a Sandbender preset before it creates anything.

## Where Presets Live

The folders and presets in the browser correspond to folders and files within Sandbender's own folder on your storage device. This folder is located at `C:\ProgramData\Audio Damage\Sandbender\Presets\` on Windows, and `~/Music/Audio Damage/Sandbender/Presets/` on macOS and Linux. For your presets to show up in the User list they must be placed in the User folder within it. Also, to avoid possible collisions during future updates, do not store your presets within the Factory folder.

Sandbender also tells hosts that support it where its presets are, so that you can browse and load them from your host's own browser without opening Sandbender's window. The host sees each preset's name, author, notes and tags.

## Default Preset

If you save a preset with the special name "Default" in the User folder, new instances of Sandbender will load it automatically when you add them to your DAW session. The standalone version opens on it every time it starts. You can use a default preset file to give yourself the same starting point with Sandbender whenever you use it.

## What Happens When a Preset Loads

Changing presets on a synthesizer that's playing can be an ugly business, as notes from the old sound collide with settings from the new one. Sandbender avoids this by fading its output away over a few milliseconds, silencing every voice, clearing the delay and reverb so that no echo of the old sound survives into the new one, setting every control to its new value at once rather than gliding there, and fading back in. The whole thing takes about a fortieth of a second. Notes you are holding are cut off, and latched notes on the keyboard are released, so play again to hear the new sound. The PERFORM surface lets go of every control its nodes were holding and takes up the new preset's axis settings, and its loop stops and is replaced by the new preset's own, or emptied if the preset doesn't have one.

The exception is the Genetics page, which is designed to change sounds under your fingers: auditioning a sound there doesn't fade or cut anything.

## And Finally...

Thanks for purchasing Sandbender. We make every effort to ensure your satisfaction with our products and want you to be happy with your purchase. Please write to support@audiodamage.com if you have any questions or comments.

## Document Revisions

- 2026-11-03: Initial draft for version 1.0