Serum

a Synthesizer Instrument plug-in
for VST, AudioUnit, and AAX hosts

Version 1.0.0
September 2014
© 2014 Xfer Records
www.xferrecords.com

**Serum Plug-in
for Windows and Macintosh OS X**

Coding: Steve Duda (www.steveduda.com)

Manual: Steve Duda

Manual version: Thursday, September 18, 2014

VST Plugin technology by Steinberg
VST is a trademark of Steinberg Media Technologies GmbH
All other copyrighted trademarks belong to their respective owners.

W W W . X F E R R E C O R D S . C O M

1. Contents

2. Introduction

Serum is a synthesizer instrument allowing for creating unique sounds in a variety of ways! Reading this manual is highly recommended, so nice job so far!

Registration

If you purchased Serum through Xfer Records, you are automatically registered for free updates.

Getting in touch

If you experience any problems while using this plug-in, or you just wish to pass on your comments regarding Serum or this manual, or links to cool tunes made with Serum, you can email the developer directly at: steve@xferrecords.com.

Installing Serum

Serum is available in both PC and Mac OS X versions. It is designed to be used within a host audio application that supports the VST plug-in format.

- ### Installing on the PC (Windows)
  Unzip the .zip file. Place Serum.dll in your VSTPlugins folder (or subfolder inside), as determined by your host program. **Place the "Serum Presets" folder in the same location as Serum.dll** (your VSTPlugins folder, or subfolder thereof).

- ### Installing on Macintosh OS X
  Unzip the .zip file. Place Serum.vst in the VST Folder of your choice (/Library/Audio/Plug-Ins/VST/). **Place the "Serum Presets" folder in /Library/Audio/Plug-Ins/VST or /Library/Audio/Presets/Xfer Records/**

- ### Uninstalling Serum
  PC: There are no registry modifications, and thus no uninstaller. Delete the Serum plug-in(s) from the folder(s) you chose on install. Delete the "Serum" folder in C:/Documents/(you)/AppData/Roaming/Xfer/ and the "Serum Presets" folder in C:/Documents/Xfer/

  OS X: remove Serum file/folders from the following locations. (command-shift-G to go to folder):

  /Library/Audio/Plug-Ins/VST/

  /Library/Audio/Plug-Ins/Components/

  /Library/Audio/Presets/Xfer Records/

  /Library/Application Support/Avid/Audio/Plug-Ins

_____

3. Getting Started

Step 1: Add Serum

Serum should be added to an Instrument (MIDI) Track.   This is the same procedure you would follow adding any software instrument in your host. If you are not sure how to do this, refer to your host documentation or contact us through the website, and we'll try to help you out.

At this point, sending MIDI notes to Serum should trigger the default sawtooth sound. Similarly, you should hear Serum's sawtooth output if you click the mouse on the piano keys at the bottom of Serum's UI.
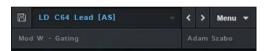
Oscillator panel, with 3d waveform display showing.

Step 2: Load a Preset

Preset Menu

Assuming you don't want to make your own sound from scratch, a good place to get started is by

Preset Menu. Clicking the main (blue) title invokes a menu, containing all of the presets in Serum's folder.

loading or browsing Presets. Select a preset by clicking on the word "- Init -" a pop-up menu will appear, displaying the available presets. Once you have selected a preset, you can click on the < > arrows to the right of the preset display name, to navigate through this preset subfolder without having to repeatedly display the presets menu.

The disk icon to the left of the preset title is for saving your own preset. In the lower half of the preset bar (pictured above) are two text fields for Preset Notes, and Preset Author and Notes.   Double-click one of these to type your own text, which will be saved when you go to save the preset.

Author and Notes text fields
Here you can type a very brief description for you and your sound. For a description, when in doubt, it is recommended to mention an important mod assignment here, such as "Mod Wheel to Chaos amount").

Step 3: Mouseover Help (Tooltips)

If you hover the mouse pointer over any control (knob, etc) for several seconds, mouseover help will appear which can give you more information about a particular control/feature without needing to visit this manual!

These are also useful for when you have assigned Modulation, for instance if you had modulated Filter Cutoff and wish to see what mod sources are controlling the knob at-a-glance, simply hover the mouse over the filter Cutoff knob and they will display, as pictured to the right.

Fil Cutoff

Assigned Modulators:
LFO 3
Env 2

Hovering the mouse pointer over a control which have modulation assigned will display the mod source(s). In this case, hovering the mouse over the Cutoff knob indicates that LFO 3 and Env 2 are assigned to modulate the Filter Cutoff.

4. Wavetable Oscillators

Wavetable oscillators are at the heart of Serum's sound generation. Unlike many wavetable synthesizers, Serum's tables are multi-cycle, which means you have a greater variety of sound.
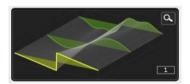
Serum's oscillator playback has been carefully constructed to give you a solid high-frequency representation all the way to the limits of the human ear, without the audible aliasing artifacts / nyquist reflection commonly found on most wavetable synthesizers. While this takes more CPU on both load-time and run-time, Serum has advanced SSE optimizations to minimize the CPU expense, and we think you'll agree that the benefits in sound purity make it worthwhile.

What is a wavetable?

A "Wavetable" is a name given to a small amount of digital audio (sample data, or waveform) which is played back in a looping fashion. The frequency (pitch) of the resulting note is created by the rate at which this waveform is played back (the speed between the repeats). The tone (harmonics) of the sound are based around the content within this waveform.

Anatomy of a Serum Wavetable

The wavetables in Serum consist of up to 256 sub-tables, or single-cycle waves which will be referred to hereafter as *frames*. This is probably best thought of as (up to) 256 discrete waveforms, which are joined together end-to-end in the parent file on disk. In normal circumstances you hear one of these 256 tables at a time, but you can automate (using the WT Pos knob) to have the sound 'animate' though various tables. Technically it is possible to hear multiple tables simultaneously, by enabling Unison (for multiple simultaneous voices) and the Unison WT Pos control (to have said multiple voices play different sub-tables, see Global settings chapter for more on this control).

Three frames (yellow = current, and two green). Notice the interpolated tables between them (faint gray).

Technically speaking, when Serum loads a wavetable it is using 2048 samples for a sub-table of the wavetable set. This means the maximum file size would be 2048(samples) x 256(sub-tables) x 32(bits) (which is exactly 2 megabytes).

Most wavetable files won't be this large however. It is common to have a good sounding wavetable which may consist of just a few sub-tables.

These sub-tables can be *interpolated tables* which can get created in the wavetable editor, to allow for smooth-sounding transitions from one to another. These interpolated tables are generated by crossfading (mix blend) or spectrally morph (frequency+phase blend, as pictured above (grey areas). These sub-tables are computed at load-time: Serum embeds the information for which interpolation type is in use, rather than the interpolated waveforms (which eliminates wasted disk space).

Wavetable Osc Module

Serum features two wavetable oscillators which are typically the primary sound synthesis source. (In addition there are the noise and sub oscillators, discussed separately later).

Osc Enable

The top-left has a power button (pictured enabled, in blue, in the diagram above). You can use this to mute an oscillator which will free up CPU. The entire panel will turn grey (shaded) to give you a clear visual that the oscillator is silent/off. In addition to CPU savings, The Osc Enable switch can be useful for temporarily *muting* an oscillator, which can be desirable when working on other oscillators, for instance.

Waveform display

The green waveform area has one of two view options – 2d and 3d. Click this area to toggle between two viewing modes. In 2d mode, you see the current single-cycle of the waveform only. The 2d view has the benefit of being able to visualize the Warp feature (described later in this chapter) in realtime. The 3d view (pictured above) allows you to see all of the sub-tables at-a-glance. Every *frame* is represented by a green horizontal waveform. Interpolated frames will be represented in gray. The currently selected waveform will draw in Yellow, and it's table number will display in the lower-right of this area (notice the [1] in the diagram above).

Waveform display magnifier

In the top-right of the Waveform display is a magnifying glass icon, clicking this opens the Wavetable Editor, which has it's own chapter in this manual.

Wavetable Menu

Clicking on the Wavetable name (displays "Default" by default / "DudaChoir3" in the diagram above) brings up a pop-up menu with factory wavetables (and user ones too, if you've created your own). Similar to the Preset menu, clicking the arrows < > allows for navigating through the wavetables in the last selected subfolder.

Wavetable Save

If you have edited the wavetable in the Wavetable Editor and/or imported your own sound, you will see the Wavetable name displaying in red, indicating that the wavetable is edited and not saved as a file on disk. You will also see a floppy-disk icon appear to the left of the table name. Clicking this disk icon will present a dialog window for saving your own folder. It is advised to save your own tables to the User folder, and very strongly advised not to overwrite the factory tables (or presets may sound different). When in doubt, always pick a new name.

> NOTE: Serum will **always** save user-made or user-edited wavetable data inside a preset (your song) unless the wavetable is Factory. While this uses hard disk space (how much size depends on how many frames you use in the Wavetables, from 8k to 4 Megabytes) the benefit is that you can exchange presets with others, or open your song in the future, without having to worry about table file management.   so there is no need to save wavetables unless you want your wavetable to appear in the Wavetable Menu.

> This does increase the size of presets and therefore your host song files as well There is a setting in Serum.cfg you can change, in order to disable storing of saved tables in to presets, but this is not recommended to change typically.

Pitch Controls

There are 3 numeric readouts for pitch in the top-right area (0 | 0 | 0 | ---- as pictured above). These are left-to-right: Octave, Semitones, and Fine (cents), and Coarse (CRS).   It was decided to have these as separate controls, rather than e.g. "36.04", for the purpose of automation and modulation – sometimes it is nice to assign an LFO to octave, for instance, or change octaves without counting in 12's.

 However, sometimes you might want to be able to have an LFO control oscillator pitch in a more coarse manner (for instance, for a siren-type sound, you might want the LFO to affect the oscillator pitch smoothly across an octave or more). To do this, drag the

Oscillator panel, with arrow pointing to the pitch controls row, described here

modulator source to the CRS section (to the right of the '0 | 0 | 0' Oct Semi Fine numbers).   You'll see a numeric box appear and you can control the mod depth directly with the mouse on this numeric box. (alternatively, like all routings, you can use the matrix to make this assignment – it is named 'Osc Coarse Pitch').

Unison

Just below the Waveform Display to the left are the unison controls (number box and 2 knobs). The number box determines how many Unison voices are used. This is a special stacking of the oscillator similar to playing multiple notes of the same pitch, but slightly detuned. Because the synthesizer is in fact generating multiple voices, this does raise the CPU demand.   The color of the number box will change as you increase the unison to serve as a reminder of the CPU consumption. The classic "magic number" for unison is 7. Serum will allow you to stack up to 16 voices for the unison, but it isn't traditional to use so much unison: as you stack more voices the resulting sound becomes more "cloudy" and less like distinct voices. As a feature, Serum gain-scales based on the unison number (the output is reduced to approximately compensate for the increase in level by the unison stacking), so you aren't fooled by the increase in volume that unison would otherwise provide.

Unison Detune – Only applicable when unison is enabled (set above "1"), this is the amount of tuning offset +/- for the additional voices.

Unison Blend – Only applicable when unison is enabled (inf act Unison must be set above "2"). This is the amount of level offset of the unison voices versus the 'central' unison voice or voices (1 if an odd number, 2 if an even number of unison). This allows you to get a "wet/dry" blend of sorts between a unison and non-unison sound. The default value of 75% is an even blend between all of the voices.

WT Pos

This sets the position within the wavetable, in other words which sub-table you will hear. In the picture on the previous page, you will notice that the 'WT Pos' knob is at the minimum setting (0%, about 7 O'Clock). This is why the first (low/left) wavetable is hi-lighted in yellow on the wavetable display. The yellow indicates this is the sub-table which will be audible.

Warp Menu

 In the lower-middle of the WT Oscillator is a menu under a knob, which displays 'Off' by default.   This menu allows you to select a Warp mode, which allows you to manipulate the playback/sound of the WT Oscillator. You can use the tiny arrows to the left of the menu to switch through different Warp Modes without having to keep opening the menu. The knob will set the depth.   If you have the waveform display in 2d view (single table visible), you will be able to see how the warp mode affects the waveform for most modes (not visible with FM or AM modes).

Below is a table of Warp modes and descriptions of their function.

Self Sync	Self-sync performed on the oscillator. A formant-shifted type of sound property. The wavetable plays/repeats faster than normal and abruptly restarts at the 'correct' length (note frequency). Best used for buzzy sounds like square/saw etc: This results in a bright/edgy sound which can sound harsh if you are using a soft-sounding waveform, which is why there is the next option (Windowed Sync) as well...

Windowed Sync	Technically should be called "windowed sync". Provides a smooth fade across the waveform so as to soften discontinuities created by the self-sync modulation.
Bend +	'Pinches' or bends the waveform inwards towards the middle of the wave cycle.
Bend -	'Pulls' or bends the waveform outward towards the edges of the wave cycle.
Bend +/-	Allows for both of the above, depending on WT Pos knob value: 12 O'Clock on the WT Pos knob (50%) represents no change to the sound.
PWM	Pushes the entire waveform to the left. This is useful on square-wave type sounds especially for the classic "PWM" sound, but useful on other waveforms as well.
Asym +	Similar to "Bend", but bends the entire waveform to the right, instead of both halves of the duty cycle separately.
Asym -	bends the entire waveform to the left.
Asym +/-	bends the entire waveform to the left or right.
Flip	Creates an instantaneous polarity flip (often called phase-inversion) on the waveform. The WT Pos determines where in the duty cycle this flip occurs.
Mirror	Creates a mirror-image of the waveform for the 2nd half of the duty cycle. Has an 'octaved' type of quality to the sound due to the 'doubling' of the waveform in to both halves of the wave cycle. WT Pos knob behaves similar to "Asym +/-", except on both halves independently, and the sound never approaches bypassed due to the mirroring of the waveform.
Remap 1	Custom remapping of the wave cycle. A magnifier glass will appear when selecting the remap feature, which opens a graph to change how the waveform will remap. A diagonal line from lower-left to top-right would indicate no change to the waveform (y=x). The WT Pos knob determines the strength of the remap, from 0 (y=x) to 100% (what you see on the graph).
Remap 2	A Mirrored remapping. Same as above, but applies the graph to each half the waveform independently. This allows for symmetric remapping, without having to attempt to draw symmetric shapes on the graph.
Remap 3	Sinusoidal remapping. Another option for remapping without attempting to draw fancy curves.
Remap 4	A 4x remapping. Similar to Remap 2 (mirror) but now the graph takes place 4 times. This creates a more busy sound, and can just be an interesting option for when you want something nasty.
Quantize	Similar to sample-and-hold, that is sample rate reduction. As

	opposed to a SR Redux **effect**, this is taking place on the waveform itself, so the aliasing sound will follow pitch perfectly (instead of having that 'same ringing pitch on all notes' quality that a redux effect gives you).
FM (from other OSC)	Frequency Modulation from the other oscillator. The other oscillator must be enabled for this to function. However, you can turn down the volume of the other oscillator if you wish to use the other oscillator simply as a modulation osc. Please note, you can only use A->B or B->A and not both- if you attempt to select this in both menus, you will notice the first assignment will turn off.
AM (from other OSC)	Same as above, except Amplitude modulation instead.
RM (from other OSC)	Same as above, except Ring modulation.
FM NoiseOsc	FM from the Noise OSC. What can I say, noise modulation can be fun :-). You need the Noise Osc to be enabled, but not audible (noise osc volume can be at zero).
FM Sub Osc	FM from the Sub Oscillator. You need the Sub Osc enabled, but the volume can be at zero.

Phase Knob – This determines where the oscillator will begin playing back when a note is triggered. It is the same concept as "sample start" on a sampler (except our "sample" is a very small waveform). Since oscillators tend to be reasonably high frequency, you may not notice a difference changing this - in particular, if the adjacent Random knob is up considerably you probably won't hear a difference adjusting the phase knob (because Random will be altering the phase on each new note). However the effect can be very pronounced in a couple scenarios:

- smooth waveform and fast-attack envelope: if your wavetable is a sine wave for instance, and the ENV1 attack is very fast (small value), then you will hear a "click" at the beginning of the sound if the Phase Knob is set to a non-zero crossing (0% or 50% represent the zero crossings in a default sine wave).

- Using two oscillators (without Random): The interaction between two oscillators can be very noticeable (phase cancellation). Adjusting the start time of one of the two oscillators will result in a different tone.

If you have the oscillator's waveform overview in 2d (single table view), as you move the Phase knob, you will see a yellow line appear indicating where the note onset will occur.

If you set the Phase knob to 100%, you will see "Mem" will appear instead of 100% (100% is the same as 0% so it is not needed). This is a special-case where the phase will use "Memory", in other words it will

continue from the location of last note.

Random Knob – This knob can be thought of as altering the Phase Knob's value by a random amount for each new voice. "Why would I want to randomize the start phase of an oscillator?" you ask. Well, this has several core uses:

1) to provide a different "start" or click/thump to each note.

2) To provide a random "tone" to each note (when layering multiple oscillators), as the phase cancellation between the oscillators will vary with each new note.

3) To reduce or remove the "Laser Zap" effect when unison (slightly detuned) notes are triggered

Here is an example if you wish to "hear" the Random Knob:

1) set Phase Knob to 0%

2) set Random to 0%

3) raise Unison # to something high, e.g. 8

4) lower Unison Detune slightly, e.g. 10 o'clock (0.10 value)

At this point you should hear a "laser zap" phasing sound stemming from the unison oscillators all starting together in phase, and slowly drifting apart from their detune. Every time you trigger the note, you will have this same laser zap sound, since the voice phases are restarting. While it can be cool sometimes, generally this is undesirable as this sweeping sound can be distracting. As you raise the Random knob, you will notice this effect become less pronounced, as a random phase offset is introduced into each voice separately. By the time you reach 100% on the Random knob, there will be no "zap" sound occurring anymore.

5. Noise Oscillator

The noise oscillator in Serum allows for more functionality beyond a typical white-noise sound. It is in fact a stereo sample player with high-quality playback.

Noise Osc Controls

Noise Menu – Similar to the Wavetable Menu, clicking the text brings up a pop-up menu with available noises (and user ones too, if they exist).   Clicking the arrows < > allows for navigating through the noise sounds in the last selected subfolder.

OneShot/Looping – this toggle button displays either an arrow pointing to the right (One-Shot). When deactivated, the sample plays indefinitely, looping start to end. When Enabled (One-Shot), the sample stops when playback reaches the end of the sound file. Looping is the default and typical-use for noise sounds, but one-shot is useful for attack sounds, for instance to add a percussive punch/attack transient to a sound.

Phase Knob – Same concept as the Phase Knob for the oscillator, but since a noise sound is a lot longer than an oscillator waveform, it might be easier to think of this as "sample start" like you would find on a sampler.   The knob can be automated/adjusted which results in a sort of lo-fi "scratching" effect.

Random Knob – Same concept as the Random Knob on the oscillators. Raising the Random will prevent the noise from being "identical" each time you press a note (holding a chord is a good example of where you would likely want to use Random, otherwise the exact noise sound will exist for all 3 notes).

Keytrack Switch – Determines whether or not the Note changes the pitch amount. When enabled, this switch puts the Pitch knob in to a different state (snaps to semitones).

Pitch Knob – Determines the base pitch/frequency for the noise oscillator. The default (50%) is nominal pitch, that is the original pitch of the sound file (assuming keytrack switch is off).

Pan Knob – the placement in the stereo field for the noise oscillator (Left-to-right).

Level Knob – The output volume of the Noise Oscillator.

Noise Oscillator sounds

The sounds of the Noise Oscillator are mono or stereo wav files. While you can put your own samples in the folder (Serum Presets/Noises). [note: you can access this folder on your hard disk easily from Serum's main menu->"Show Serum Presets Folder"].

If you choose to place your own samples in the Serum Noises folder, there is one important consideration:

<u>Noise sample data does not get copied in to the patch/preset.</u>

This means if you open a song in the future (say, on a different computer in the future) you will need to make sure the Noise sound is located in the same folder within the "Serum Presets/Noises/" folder. Similarly, if you make a patch and want to share it with a friend who also owns Serum, they too will need this noise sample. If the noise is not present, Serum will alert with a dialog, showing the missing sample location. Place the noise sample in the specified location and reload the preset/song.

The noise oscillator playback is a high-quality realtime interpolation. This is because being noise sounds, the high-frequencies matter, and it was a design decision to use the best quality wherever possible, even at the expense of additional CPU (although it is heavily optimized with both SSE2 and pre-calculations when a file loads).

Since this high-quality playback does take CPU, it is worth noting that mono sounds do take slightly less CPU (since it is one channel instead of two, and this can add up with chords, etc). The main consideration here is if you have a mono noise source, you shouldn't export it as a stereo file for Serum, or you will be wasting not just disk space, but CPU on playback.

6. Filter Module

The Filter module gives you per-voice filtering of one or several oscillators. The buttons in the Lower-left "**A**", "**B**", "**N**", and "**S**" route the corresponding oscillators to the filter (osc **A**, osc **B**, **N**oise osc, **S**ub osc). In the top-left is an **Enable-switch** for the Filter module. This prevents an easy way to disable the filter altogether.

Filter Type
Filter Type is chosen with the menu at the top of the module (pictured to the right, says "MG Low 12" in the picture). Clicking this text brings up a pop-up menu for selecting the filter type. Alternatively you can use the < > arrows to the right of the name for advancing through the filter types with a single mouse-click. See the appendix at the very end of this manual for a list of the filter types.

Cutoff
knob sets the primary cutoff frequency for the filter, with just a couple exceptions (such as vowel for the formant filters). The key-track switch (piano icon, in the very lower-left of the module) allows the cutoff to be offset from MIDI notes (in most filter types, one octave of MIDI is precisely one octave of filter frequency control).

Res
The Res knob sets the resonance, or feedback of the filter circuit.

Stereo
The Stereo knob creates a cutoff *offset* for the left and right signals. At 50% (12 O'Clock, default) this knob has no effect. When turned to the left (counter-clockwise) the Left channel cutoff will increase, and the Right channel cutoff will decrease. When turned to the right past 12 O'Clock (clockwise) the opposite will happen: the Left cutoff will decrease and the Right channel cutoff will increase.

Drive
The Drive knob increases the gain into the filter circuit and can impart some coloration (mild distortion) on to the sound.

Var
The Var knob (says "Fat" in the picture above) is a Variable knob, which has a different function depending on the filter type chosen. For

instance in the "dual" filters, this knob is a control for the second filter's cutoff. See the appendix at the very end of this manual for a list of the var parameter functions.

Mix

The Mix knob allows for a wet/dry amount for the filter. The default (100%) means 100% wet.

Bonus: You can click on the filter display to see the filter phase response! Probably more eye-candy than useful..

7. Effects

Serum offers an Effects page consisting of 10 effects modules. These modules can be re-ordered and you are able to use up to all 10 at once.

To Access the Effects page, click "FX" at the top of the window. You will now see the rack and the 10 effects names down the left side. To enable an effect, click the corresponding switch along the very left edge of the window. You will now see the LED illuminate as well as the effect module will appear to the right in the Rack.

All audio output from the oscillators+filters route to the FX module (and then to the master volume, and output.

The signal flow is top-to-bottom through the FX. To re-order the effect order, drag these name tiles (notice the up+down arrows icon) up or down in the 'list'. The effect modules in the Rack always draw in the the same order as the tiles on the left.

Common controls to each FX module

MIX: Every effect (except EQ) has a Mix knob on the right-most knob of the effect. This knob sets the wet/dry balance of the signal (knob at 0 = 100% dry, knob at full = 100% wet).

Bypass: At the very right edge of the FX module is a power button. This is a bypass, intended really for "temporary" use for example to hear your sound with and without a given FX (without having to resort to setting the MIX knob to 100% dry). To "permanently" deactivate an effect it is recommended to use the very left-most enable buttons, because these will also hide the module from view.

Modulating FX parameters

Similar to modulating the standard synth controls (see next chapter) you can also modulate most FX parameters. It is worth noting however that the FX rack is a DSP process which takes place on the "sum output" of the synth engine, rather than per-voice. You could think of this as "The effects are monophonic" or "the effects are like plugin inserts after Serum". Therefore when playing polyphonic synth parts (strummed chord, for example) you should keep in mind that automating FX controls with per-voice mod sources (such as an envelope) will result in the effect parameter modulation will be getting modulated/re-triggered

by each new note.

Individual Effects modules explained

Reverb
The reverb is a Plate reverb algorithm courtesy of Togu Audio Line. It is a modified version of the Tal Reverb algorithm:
http://kunz.corrupt.ch/products/tal-reverb

Size: Sets the room size (reverb time + dimension)

Pre Delay: The amount of time in milliseconds before the reverberation occurs. Using pre-delay allows you to give the impression that a sound is close to you but in a large room, separate your transient from the reverb, or create a delay-like echo.

Low Cut: suppress low frequencies from the reverb. 0% = no effect on lows, 100% = no lows at all

Damp: an additional high-frequency cut for the reverb. The knob controls how fast this high-frequency attenuation occurs. 0% = no damp, 100% = maximum damping.

Hi Cut: suppress high frequencies from the reverb. 0% = no effect on highs, 100% = no high frequencies at all

Width: expand or collapse the stereo width of the reverb. 100% = maximum width.

EQ
The EQ offers two-band parametric control.   Two 3-state switches in the middle of the EQ section (adjacent to/touching the center screen) allow you to set the type for each of the two bands. The left band is intended for low-frequency (LF) adjustment and allows for Low Shelf, Peaking, or hipass filtering. The right band is intended for high-frequency (HF) adjustment and allows for High Shelf, Peaking, or Lowpass filtering.

The knobs act as follows:

Freq (L): set the frequency in Hz for the low EQ band.

Q (L): set the Q (resonance) for the low EQ band.

Gain (L): set the gain boost/cut in dB for the low EQ band.   This knob has no effect if you have selected Hipass as your low EQ band type.

Freq (R): set the frequency in Hz for the high EQ band.

Q (R): set the Q (resonance) for the high EQ band.

Gain (R): set the gain boost/cut in dB for the high EQ band.   This knob

has no effect if you have selected Lowpass as your high EQ band type.

Distortion

The Distortion module allows for 13 different distortion types, including 2 dual-waveshaper modes which allow you to create your own custom distortion. The controls from left-to-right are as follows.

Mode: The menu (displays "Tube" by default) brings up a list of the distortion types. The < > arrows allow you to advance through the modes without having to click the menu.

Off/Pre/Post Switch: This switch enables the filtering, which can be set to either Pre-distortion or Post-distortion.   When enabled you'll notice the filter display will turn orange, indicating that the filter is active.

F (frequency): This number-box under the filter display sets the cutoff frequency for the filter. [requires filter enabled, see previous control]

Q (resonance): Sets the resonance for the filter. Allows for high values, for squelchy feedback. Typical use is a lower setting.

LP/BP/HP: Sets the filter type for the distortion filter. Allows for a morph between lowpass-to-bandpass, or bandpass-to-hipass.

Drive: Essentially the "Amount" control: typically the amount of gain boost for the distortion. The exceptions are: on "Downsample" effect type, the Drive knob controls the amount of sample rate reduction.   On the "X-Shaper" and "X-Shaper (Asym)" modes (read more below), the drive knob is a 'morph' between the the two waveshapes.

X-Shaper (dual waveshaper) FX modes

The X-Shaper is a "dual crossfading waveshaper". The Drive knob, as mentioned above, controls the blend between the two waveshaping graphs (0% drive knob = waveshaper A, 100% drive knob = waveshaper B).   When you have selected an X-Shaper from the distortion type menu, you will see "Edit A" and "Edit B" buttons appear below the distortion type. Clicking either of these will bring up a pop-up editor for the wave. This editor is an "X-Y" graph, with X (horizontal) representing input level, and Y (vertical) representing the corresponding remapped output level for any given X input level.

"**X-shaper**" is a symmetric waveshaper, the lower-left point on the graph represents silence (-INF dB for input and output), and the top-right point representing the highest level (0 dB for input and output).

"**X-shaper (Asym)**" is an asymmetric waveshaper. In this case, the **middle** of the graph represents silence (-INF dB input and output), the top-right represents the highest positive value to the signal, and the lower-left represents the highest possible **negative** value.   Asymmetric distortion allows you to bring out the even-order harmonics not typically found in a standard symmetric distortion (such as clipping). This is often the case in guitar amps, one pole will be distorting (e.g. flatline) while the other pole may remain relatively undistorted.

Flanger

BPM Sync: This switch determines whether or not the flanger's sweep is synced to the host BPM. The rate knob will now snap to musical times. When the switch is off, the rate knob is in Hz.

Rate: As mentioned above, the BPM switch determines the Rate knob's exact operation- either between 0 Hz and 20 Hz (BPM Sync off) or 0 Hz and 1/32nd notes (when BPM Sync is enabled).

Depth: How much the Flanger's LFO influences the sound, in other words how much/deep the flange is taking place.

Feed: The feedback amount of the flanger circuit, makes the effect more pronounced ("ringing").

Phase: The stereo phase offset for the LFO influence over the flanger (Left flange and Right flange offset).   0% = both L and R have the same frequency. 50% = 180 degrees - L and R have opposite frequency, in other words the flanger sweep will be "rising" on the left while "falling" on the right, or vice versa.

Phaser

BPM Sync: This switch determines whether or not the phaser's sweep is synced to the host BPM. The rate knob will now snap to musical times. When the switch is off, the rate knob is in Hz.

Rate: As mentioned above, the BPM switch determines the Rate knob's exact operation- either between 0 Hz and 20 Hz (BPM Sync off) or 0 Hz and 1/32nd notes (when BPM Sync is enabled).

Depth: How much the Phaser's LFO influences the sound, in other words how much/deep the phase effect is taking place.

Freq: Sets the base frequency for the phaser effect.

Feed: The feedback amount of the phaser circuit, makes the effect more pronounced ("ringing").

Phase: The stereo phase offset for the LFO influence over the phaser (Left flange and Right flange offset).   0% = both L and R have the same frequency. 50% = 180 degrees - L and R have opposite frequency, in other words the phaser sweep will be "rising" on the left while "falling" on the right, or vice versa.

Chorus

The chorus module is a 4-voice chorus effect, with 2 L and 2 R chorus taps.

BPM Sync: This switch determines whether or not the Chorus's modulation rate is synced to the host BPM. The rate knob will now snap to musical times. When the switch is off, the rate knob is in Hz.

Rate: As mentioned above, the BPM switch determines the Rate knob's exact operation- either between 0 Hz and 20 Hz (BPM Sync off) or 0 Hz and 1/32nd notes (when BPM Sync is enabled).

Delay 1: Sets the amount of delay time in milliseconds between the dry signal and the first stereo pair of chorus voices.

Delay 2: Sets the amount of delay time in milliseconds between the dry signal and the second pair of chorus voices.

Depth: sets how much the delay times above are modulated by the Chorus's LFO. In other words, how much pitch warble occurs.

Feed: the feedback amount of the chorus voices, in other words how much of the chorus voice's output appears back at the input of the chorus module. This creates a more pronounced "ringing" to the chorus.

LPF: set the cutoff frequency in HZ of the lowpass filter which is after the chorus wet effect. This is useful for a more (or less) "warm doubling" of the signal.

Delay

Feed: sets the amount of the delayed signal appearing back at the input of the delay, useful for controlling how many repeats of the delay are audible.

BPM Sync Switch: This determines whether or not the delay times are in tempo-based units (¼ note, etc). Or in milliseconds.

Link Switch: When enabled, the right-channel delay times will be linked to (kept the same as) the left channel delay times.

Delay times: You'll notice two columns of number boxes, labelled "LEFT" and "RIGHT". The upper box is the base Delay time for the respective Left or Right audio channel. The lower box is an Offset (scalar) for the above delay time, for instance a value of 1.1 means that the above delay time will now sound 110% of its value. You'll also notice when dragging these lower boxes you'll see "TRIP" and "DOT" appear when you reach 133% (1.333) and 150% (1.5) respectively. These also act a little "sticky" with the mouse movement, so they are easier to find/set. This is the way you access triplet or dotted values for the delay times.

Freq (F): This number box sets the filter cutoff frequency in Hz for the delay's filter.

Q (Q): This sets the amount of bandwidth for the delay's filter (how much lowpass and hipass are applied. Technically this is working opposite a standard Q control, where typically a larger Q value means a narrow filter bandwidth. However in this case, a maximum value means minimum filtering / maximum bandwidth.

Mode Select (Normal / Ping-Pong / Tap->Delay): Normal is a

standard stereo delay with independent L and R channels/times. In **Ping-Pong**, the output of the L delay and the output of the R delay feed in to one another. In **Tap->Delay**, both delays are put in mono, in series: the L signal fires once with no feedback (tap) followed by the R signal operating as a typical delay (feedback is applied here).

Compressor

Thresh: Sets the threshold in dB for the compression to start engaging. 0 = 0 dB = no compression (unless overloaded input signal). 100% = -120 dB = almost always compressing. Typical use would likely be somewhere well in the middle of the range, eg. around -12 dB, but it is more dependent on your input signal strength, and how much compression is desired.

Ratio: sets the strength of the gain reduction. Typical compression is between 2:1 and 4:1. If you set the compression knob to maximum, you will see "Limit" appear. This is a completely different DSP circuit at that point (true peak limiter, no longer a compressor) and therefore the other controls will behave different when this Limiter is engaged.

Attack: the amount of time in milliseconds for the gain reduction to engage. A longer attack ("slow attack") is useful for letting some signal through before the gain reduction takes place, resulting in a "punch", "snap" or "bite" (or other words in your preferred language). A shorter / fast attack will tame peaks more completely.

Release: The amount of time for the gain reduction to get removed.

Gain: The amount of makeup gain. This is a good way to boost quiet signals. Allows for ~36 dB of boost – so be careful, as that is a large amount of gain! A little can go a long way.

Multiband: Enabling this switch makes the compressor become a multiband up/downwards compressor. This is an extreme setting but you may find a use for it. The individual bands are not user-adjustable separately.

Filter

The effect filter operates identical to the per-voice synth filter found on the main OSC tab, except running as a master effect. Please refer to the chapter titled "Filter Module" for specifics about the controls.

Hyper

The Hyper effect is a micro-delay chorus effect with a variable number of voices (1-7). In addition the Hyper effect can be configured to re-trigger on every MIDI note, which adds to the potential simulation of a unison.

Rate: The speed at which the various Hyper voices are oscillating sharp/flat in pitch.

Detune: Sets the amount / depth for the Hyper voice oscillations

sharp/flat in pitch.

Retrig: When enabled, the Retrig switch will reset all of the Hyper voices to start over from a zeroed pitch offset. This provides a "laser" like zap effect on each note on. Typical use probably doesn't involve enabling this switch, but you may decide otherwise, in particular on monophonic patches it may be of use.

Unison: This number box sets the number of chorus voices make up the Hyper effect.

Mix: The Hyper effect has its own Mix control, independent to the Dimension controls listed below. If you only wish to use the Dimension effect and not the Hyper, it is recommended to set the "Unison" control to 0.

Dimension

The Dimension effect is a pseudo-stereo effect made out of 4 delay lines summed out-of-phase and slowly amplitude modulated to provide a subtle amount motion to the effect. This is useful for adding a perceived of width to an otherwise mono signal.

Size: Sets the amount of delay time for the delays.

Mix: The Dimension has its own Mix control. When set to 0%, the Dimension effect is disabled.

8. Modulation Routing

Serum offers 16 Modulation Matrix slots in a patch, which allows for up to 16 destination parameters to be modified and/or scaled by up to 32 modulation sources.

Drag-and-drop Routing / Assignments

Serum contains a drag-and-drop method of routing for common connections. This allows you to establish modulation routing quickly, without the need to tediously hunt through menus, or even leave the main panel display.

This Mod Source has one destination assigned. Hovering the mouse on this tile will list all connections.

For instance, to assign LFO1 to OSCA Pan, click+drag the Tile name "LFO 1" (in the middle of Serum) to the pan knob in the Osc-A panel located nearby (diagonally up to the left about an inch from the "LFO 1" tile. As you are dragging, you should notice the "LFO 1" tile is hovering with the mouse cursor, and the mouse cursor will get a "+" sign as you hover over the pan knob. This "+" is indicating that you are over a valid mod destination. When you let go of the mouse, the connection will automatically be made, and the LFO1 is now affecting OSC-A's stereo pan position.

You'll also notice a number "1" will appear on the right side of the "LFO 1" tile. This is indicating that the LFO 1 has one destination.

Hover mouse to see where a mod source is used

You can hover the mouse over the LFO 1 tile, and it will display the assigned destination(s) ("A Pan" in this example).

Modulator depth control

Modulator depth at 100% ('halo arc').

A blue halo will appear around the Pan knob. This is indicating the LFO1 Depth on the pan knob (in other words how much influence the LFO has over the panning position). An Up/Down arrow will also appear in blue to the left of the knob. You can click+drag on this blue arrow to change the modulation depth amount (tip: you can also alt-click+drag on the knob to have the same function as clicking on this

arrow, the arrow simply saves you from having to touch the computer keyboard).

The miniature arc in the top-left of the knob allows you to adjust this depth amount by clicking the mouse and dragging.   You will notice a grey arc as well on controls which have a modulator assigned, but not the currently selected mod source. To adjust the modulator depth in that case, you must either A) first click the associated mod source tile, or B) visit the Mod Matrix window and adjust the depth from there.

Negativity (well, thats not very PLUR)....

If you lower this modulation amount beyond zero, you will notice the halo around the knob becomes a slightly different hue. This color is simply indicating that there is a negative depth amount, in other words the value is inverted (as the LFO output goes up, the influence to the pan knob goes down).

Context-Menu on controls

If you CTRL-Click (OSX) / Right-click (Windows) on a control such as a knob, a context-menu will appear. With the following choices:

• **Bypass Modulator:** bypass the current modulation connection (in other words, the connection between the selected ModSource tile, and this knob). The menu will also display this selected ModSource, for instance if LFO1 was last clicked, and LFO1 was dragged to this knob previously, you will see, "Bypass Modulator: LFO1". If no such connection exists between the selected ModSource and this control, you will see "Bypass Modulator" greyed out.   After making this menu selection, you will see the 'halo' around the knob turn grey, indicating that the modulation connection is bypassed. You can un-bypass this connection by selecting the menu item again (uncheck).

• **Remove Modulator:** Similar to above, except it removes the connection between the modulation source and this knob completely.

• **Remove All Modulators:** Removes all connections to this knob from all Modulation Sources.

• **Reset Control:** Resets this knob/control to its default value – this is the same as command-click (OSX) or CTRL-click (windows).

• **Mod Source (submenu):** Displays all of the possible modulation sources to control this parameter. This allows another way to make a connection to a control without having to drag from the

ModSource tile or visiting the Matrix window.

     • **MIDI Learn:** Selecting this menu item puts Serum in a "MIDI Learn" Mode, where it waits for an incoming MIDI CC value. Once it receives a MIDI CC, the MIDI Learn is deactivated and the CC# is now assigned to this control. This assignment is saved with the preset/patch. Revisiting this pop-up context menu again will display the CC# and give you the ability to either re-learn it to a different CC#, or else remove it.

Mod Matrix Window

Many software synthesizers including Serum have a modulation matrix. This is a list view, which allows you to select the routing and amounts for the various modulation connections. Unique to Serum is the fact that the Mod Matrix and drag-and-drop style of modulation routing both exist. Drag a mod source to a knob, and this routing will appear in the matrix, or vice-versa. This gives you additional flexibility in viewing/creating/altering mod assignments.

Mod Matrix: Sources
Some mod sources only exist in the Mod Matrix (no 'drag source' tile):

     • Aftertouch (aka Channel Pressure)

     • Chaos 1 and 2 (Chaos controls are on the Global tab).

     • Note-On Random 1 and 2 (two separate random numbers generated on a note on, in case you have a need for 2 different random values on each note on)

     • Noise Osc (the Noise Oscillator's output can be used as a modulation source, try pitching the Noise Osc low for more of a chaos LFO type of effect)

     • Pitch Bend

     • Fixed (not really a mod source, but you can use this to allow a modulation assignment to get a 'fixed' value with slider depth control, should you want that for some reason).

only exist in the Mod Matrix Source menu ("Src"). Therefore you will use the MATRIX window if you wish to use these mod sources. Furthermore the MATRIX window is used for assigning the secondary source ("Mod Src") and/or it's Mod type ("Mod") and its amount (Mod Amt).

Mod Matrix: Amount Slider
The Amount slider is bi-directional. As mentioned in the above heading

"Negativity", the slider to the left of middle will have a Negative amount value, in other words the mod source's output will be inverted before heading to its destination.

As with other controls Serum, Command-Click (OSX) / CTRL-Click (Windows) will reset the value to default (Zero in this case).

Mod Matrix: Destinations

A pop-up menu for selecting the *mod destination* (the parameter which you wish to have effected by the mod source).

Mod Matrix: Type

This determines if the modulation is <u>unidirectional</u> (displayed as an arrow pointing to the right: ->) or <u>bi-directional</u> (displayed as arrows pointing both left and right: <->). You can achieve similar sonic results with either, it depends on whether you prefer the destination control (think "Knob position") be at the beginning (Unidirectional) or the middle (bi-directional).   When creating modulations with the drag-and-drop method mentioned at the beginning of this chapter, the default Type assigned depends on whether or not the control is centered. For instance, if you drag to a pan knob which is centered, Serum assumes you wish your Mod Source to pan both left and right, so bidirectional is chosen.   If you drag to a control which is not centered, such as Filter Resonance, Serum assumes you want the modulation to **add** value only, and Unidirectional is used.

Tip: You can change the Type between Uni/Bi without visiting the Mod Matrix window, by shift-alt-clicking on the knob with a visible (blue/yellow) modulation assignment.
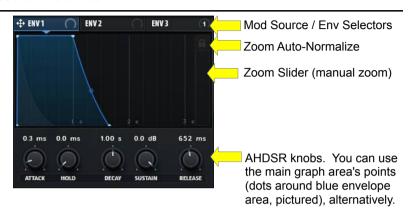
Mod Matrix: Mod

The Mod drop-down gives you some "bonus" ways to manipulate modulation assignments.

*	Default. The two modulation sources are "multiplied" together, so that one is scaling the other. For instance, if LFO1 is the Src, and ModWheel is the Aux Src, then we will not hear the LFO1 influence unless the ModWheel is raised above zero (and the Aux Amt slider is also above zero).
*(inv)	Same as the above, except the secondary source ("Aux Src") is value-inverted. For instance in the above example, there would be no modulation if the ModWheel was at maximum, and there would be full modulation if the ModWheel is at minimum.
bypass	This row of the modulation matrix now has no effect.

9. Envelopes

        The Envelope module (lower-left of display) provides 3 modulation sources.

Mod Source / Env Selectors

Zoom Auto-Normalize

Zoom Slider (manual zoom)

AHDSR knobs. You can use the main graph area's points (dots around blue envelope area, pictured), alternatively.

There are 3 Envelopes in Serum.

The first envelope, ENV1, is "special" as it is dedicated for Amp (output volume of each voice), however you can also connect it to additional parameters if desired.   Clicking the names "Env 1", "Env2", "Env3" along the top of the envelope will choose which envelope to display/edit . You will still see the other two envelope shapes in the background, in grey (if they are in use / assigned to a destination).

Assigning an Envelope to a parameter

In order for an Env (Env2 or Env3) to activate, you need to first drag it's modulation source tile to a destination parameter. [note: Env1 is always active because it always controls the volume of the note]. To do this, drag the "Env 2" tile to a knob, such as Cutoff on the filter (be sure to enable the Filter module at the top of Filter).   Now lower the sustain knob below the envelope in lower-left and press a note. You should hear the sound become more filtered as the note is held down.

Env View modes
There are two different zoom display options for the envelope. A lock

button in the top-right corner of the Env determines which mode:

- – Locked: The Envelope always fills the display area. In other words zoom is set automatically so the envelope will fit perfectly in the display. In other words adjusting envelope times will simply change how much time is represented (you'll see the background time ruler change scaling).

- – Unlocked: dragging in the zoom area below the lock will determine the zoom level for the envelope display.

10.LFOs

The LFO module might look familiar to those of you familiar with the LFOTool plugin! There are in fact a lot of similarities, and some extra features as well.

There are 4 LFO's which each contain an entire set of independent controls. You select which LFOTool to view or edit by clicking the Modulator Source tiles ("LFO 1", "LFO 2", "LFO 3", "LFO 4") directly above the LFO Graph.

The LFO Module has its own set of controls (pictured) below the (blue) shape. The settings are independent for each LFO. Only the grid size is 'global' to all LFO's.

LFO Graph Display

The graph area is where you draw the shape for the LFO. Here are the mouse/hotkey functions available for the graph:

• **double-click** to Add or Remove points

• **shift-click** to draw steps at the Grid Size (step sequencer)

• **alt-click+drag Points** to snap points to the Grid Size

• **alt-click+drag any Curve Point** to move all curve points at once.

• **click+drag on background** to multi-select points

• **command-click+drag a point (ctrl-click on Windows)** to multi-select points for a relative movement (rainbow color will appear on points, dragging them in this state will make closer points move more and further points will move less).

• **control-click (right-click on Windows)** to bring up a pop-up menu for additional features, such as setting the segment 'shape' for shift-click, deleting all multi-selected points, or assigning the start or loopback points.

• **shift-command click (shift-control-click on Windows)** on a point when in Env mode, to set this point as the Loopback position (or the very last point if you wish no Loopback position). This is simply a shortcut to avoid the menu, and probably no need to remember this one

since you can use the aforementioned pop-up menu to access this feature.

The LFO Controls

The LFO module is found at all times in the bottom-center of Serum window. It has its own set of controls as well, from left-to-right:

Grid Size
 This number box allows to set the grid size of the LFO graph. You will see the visual grid background of the LFO graph change while adjusting this number. The Grid exists to be used in conjunction with the Alt-click (snap points) or the Shift-clicking (draw step segments) on the LFO graph.

Load/Save menu
 This folder icon will display a pop-up menu when clicked. There are many presets included which will overwrite the visible LFO graph. The bottom option of this menu is "Save LFO Shape.." which will allow you to save your own LFO Shape for future use.

LFO Mode (Trig / Env / Off)
 Allows you to select how the LFO behaves to new notes:

   • **Trig**: re-triggers the LFO. the LFO will start with a new note. This is useful for many situations when you want the LFO to always behave the same timing-wise with new notes.

   • **Env**: similar to Trig, except the LFO will stop once it reaches the right-edge of the graph (plays once, like an envelope). It is possible to loop a segment of the LFO, for instance add a point to the middle of the LFOGraph and shift-command-click on it (shift-ctrl-click on windows) and you'll see an orange marker appear above the point. Now when you play a new note and hold it down, assuming you have routed this LFO to a destination, you'll see+hear it looping this "second half" of the LFO graph (from the marker to the right edge).

   • **Off:** Note on's are essentially ignored, the LFO will run mindlessly on its own like an analogue LFO on a synthesizer would. This is useful for creating a more free-form modulation which ignores new notes, and/or one which is always synced to your song's time position.

BPM switch
 Enabling this will make the time value snap to tempo-based units (¼ note, 1/8th note, etc).

Dot/Trip switches
 (for when BPM switch is enabled). These allow for triplet and dotted times on the rate control, respectively. These switches are useful for

avoiding triplet or dotted times when automating the LFO Rate, that is, for avoiding doted and triplet times when you know you want an evenly beat-divisible time(s).

Anchor switch

(for when BPM switch is enabled). This determines whether or not the LFO's playback position will "jump" if you automate the time control: if Anchor is enabled, then the time is "anchored" to your song time, in other words changing the tempo from ¼ note to 1 bar means the time position may "jump" in order to have the playback properly affixed to the bar cycle. Conversely, if the Anchor is disabled, the play position will not jump when you change the LFO Rate, but now may feel rhythmically off-the-beat. This is how a traditional LFO behaves.

Voicing controls in the lower-right corner

Rate

the playback speed of the LFO. Said another way, this determines the amount of time represented on the LFO graph area. The LFO Rate is in beat-synced units by default (BPM switch on) but can also be made free to a Hz value (BPM switch off).

Rise

this is the amount of time taken for the LFO graph shape to have influence over the LFO output. In other words, the LFO will begin with a "fixed" output (imagine the LFOGraph is a flat horizontal line, with whatever value the left-most point of the LFOTool graph contains) and slowly (based on the Rise time) becoming the shape of the visible graph.   This knob is useful for getting the LFO to slowly have influence over your sound.

Delay

this is the amount of time before the rise begins. The LFO will have a fixed output as mentioned above, and after the delay time period, the rise will begin.

Smooth

this smooths the LFO output. This is useful for avoiding abrupt jumps in the LFO output, without having to draw ramps on every segment of the LFO graph.

Protip:

• You can alt-click+drag one of the "LFO 1", "LFO 2", etc tiles to another, to copy all LFO settings from one LFO to another.

11.Voicing & Portamento

Voicing Settings

In the lower-right corner of Serum you'll find the Voicing section. These controls allow you to change how Serum behaves to multiple notes playing at once.

Portamento controls, found in the lower-right corner (to the right of the piano keys).

Mono (Monophonic) Switch

The Mono switch enables Monophonic mode, where Serum will only allow one active note at-a-time. If a new note is pressed while one is playing, the old one is "interrupted", technically it is re-pitched to the new pitch.

Legato Switch

Legato is only audible with Mono switch is enabled. When a mono voice is interrupted, the state of the Legato switch determines whether or not the Envelopes/LFO's will re-trigger. When Legato is enabled, the envelopes do not re-trigger, which allows for a smooth change to the new note which is pressed. However sometimes you want the envelopes to re-trigger so each note has the same definition, in these situations you'd want Legato to be set to off.

Poly (Polyphony) number box

This box sets the number of simultaneous notes which can be played. Sometimes, often for CPU reasons, you would want to place a limit on the simultaneous notes. For instance if you were sending a flurry of notes (arpeggio) to a patch which had a long release (say 10 seconds) this could end up being a whole lot of voices, were there no limit (and your CPU would probably overload). Typically speaking 8 is enough and 16 is generally a bunch.   This box is greyed out (as pictured) when the Mono switch is enabled.

Poly (Polyphony) count display

Underneath this number-box is a numeric readout. In the picture above you'll notice that it says "0 / 1" in this display. This means 0 out of a total possible number of 1 voices is playing. If I held a note it would display 1 / 1.   If I enabled a 2nd oscillator, it would display "0/2" when no notes are playing. In other words, Serum totals the number of voices you have active in a patch (up to 34 synth voices can make up a single note, if both oscillators are set to 16 unison and sub and noise oscs are also

enabled).

Portamento Settings

Below the voicing settings (in the very lower-right corner of Serum) are the Portamento controls.

Portamento allows for a slow glide/pitch bend from one note played to another. It is most commonly used/most useful when Mono is enabled as well. This way if you hold one note and play another, it will slowly change the pitch from the first note to the second.

Porta (Porta time) Knob
The Porta time knob controls the rate of glide from one note to another.

Porta Curve Control
The Porta curve, adjusts the contour of glide from one note to another. If set convex (typical use), as picture above, the note pitch will depart the beginning pitch quickly and slow down as it nears the note destination frequency. If set concave (dragged down below half), the opposite is true: the pitch will slowly depart the source pitch, and later rapidly arrive at the destination pitch.

'Always' switch
When activated, the portamento will occur on a new note even if no note is currently playing.   When Off, a note must be held for portamento to take place on the (2nd) note.

'Scaled' switch
This is potentially useful for melodic leads where you wish there to be a less noticeable portamento on short intervals.

When activated, the portamento rate is adjusted based on the distance of transversal between the source and destination pitches. In other words, if the portamento is a glide between two notes 1-octave apart, the rate knob's time value will be used. If the portamento is a glide between notes **less** than one-octave apart, the time will be faster, and vice-versa (notes larger than an octave will be progressively slower).

12.Global page settings

All of the Global settings are saved in your Preset/Song/Patch, except the Preferences which are truly "Global".

Oscillator Settings

In the top-left of the Global page is the Oscillator Settings:

Pitch Tracking OSC A switch

This switch determines whether or not OSC A will follow the MIDI note to determine pitch. When off, all incoming MIDI notes will trigger OSC A at the same (low) pitch. If you wish to make this a different pitch, you can use the OCT | SEMI | FINE and/or COARSE controls for OSC A on the main OSC tab.

Pitch Tracking OSC B switch

This switch determines whether or not OSC B will follow the MIDI note to determine pitch. When off, all incoming MIDI notes will trigger OSC B at the same (low) pitch. If you wish to make this a different pitch, you can use the OCT | SEMI | FINE and/or COARSE controls for OSC B on the main OSC tab.

Noise Fine

This allows an adjustment in cents for the Noise oscillator. This is useful when Key-track is enabled for the Noise oscillator (noise pitch knob snaps to semitones in this mode), and the Noise oscillator has a discernible pitch which sounds out of tune.

Oversampling

This sets the quality for Serum. This setting applies to situations using a Warp mode (OSC Sync, FM, etc) for an oscillator. Otherwise the standard oscillator quality is always maximum!

Chaos 1 | Chaos 2

Chaos 1 and 2 are different types of chaos LFO (low frequency oscillators). To hear one or both, you must make the assignments in the Mod matrix (or by ctrl-clicking (OSX) or right-clicking (Windows) on a desired destination control (e.g. filter cutoff) and selecting the desired

Chaos under the Mod Source submenu).

Chaos 1 has a more periodic (repeating) and more bipolar (two separate tendencies, up or down). Chaos 2 has more of a chaotically 'restarting' sort of sound to it.

A good use for the Chaos modulators is **subtlety**, for instance very small fine-tune deviations to simulate the drift of analog oscillator pitch. However you may of course find your own uses!

Chaos BPM sync
Although by nature chaos is not predictable, since both Chaos oscillators do have a sort of natural frequency in the algorithms, the Chaos BPM sync will snap the LFO rates to musical divisions based on your host BPM. However you should not expect "rhythms" out, but nonetheless can be useful at the very least so you can get the sort of chaos rate you had in mind (if you think in musical units!)

Chaos Mono switch
When disabled, every voice (if you're playing a chord) will have it's own chaos oscillators.   However sometimes you may wish all voices to share the same chaos motion (for instance if controlling filter cutoff, you may want all notes to have this same chaos movement). Enabling the mono switch will make all synth voices use a single Chaos oscillator for this purpose.

Unison settings
On the global tab you can find the more advanced Unison settings for OSC A and OSC B, respectively:

(Unison) Range (maximum detune amount)
Sets the range in semitones that the Oscillator's Unison Detune knob (on OSC panel) will cover between 0 and 48 semitones. The default value is 2 semitones.

(Unison) Width (stereo spread)
Sets the stereo panning for the individual unison voices. 0 = all voices panned center, 100 = voices panned as much as 100% Left and Right.

(Unison) Warp
Sets the amount of "warp" offset for each of the individual unison voices. To hear this, you must have a Warp mode is selected (Sync, FM, etc) and of course Unison raised above 1.   An easy way to think about this is imagine the Warp knob on the Osc page has a little lower or higher value for each unison voice. The more this Warp offset is raised, the wider the spread for each unison-voice this +/- assignment, based from the Warp knob's current value.

(Unison) WT Pos
Similar concept to Warp (above) but instead the offset per-unison-voice is made to the WT Pos. This way a single note with unison enabled can be playing multiple sub-tables in your wavetable simultaneously. The

greater this control's value, the more spread.

(Unison) Stack

The stack control provides multiple options for having the unison voices playing different pitches. The default is "off" where all voices play the incoming pitch (note you played). The first option: "12 (1x)" means that every 2nd unison voice will play up an octave (0, +12, 0, +12, 0 +12, etc). The next option: "12 (2x)" means that this transposition can be either 1 or 2 octaves (0, +12, 0 +24, 0 +12, etc). The choices with +7 means that a perfect fifth will also be included in these layerings. The last choices: Center-12 and Center-24 means the center voices of the Unison stack will be lowered in pitch 1 or 2 octaves, respectively.

(Unison) Tuning

Determines the mode of detune for the detune knob. These changes are also reflected *visually* when adjusting the detune knob on the main oscillator panel, you'll notice the vertical stalks which draw on the waveform area will have a different look, depending on this Tuning setting.

"**Linear**" which means all voices will be detuned by equally spaced amounts.

"**Super**" is a little less even than Linear, the detune has a slight emphasis towards notes being the "correct" (in-tune) pitch. This simulates a little closer how an ensemble of instruments playing the same note might sound (a couple people are out of tune, but the average is closer to in-tune).

"**Exp.**", short for Exponential, favors notes closer to being in-tune even more. This might be closer to how a more talented ensemble trying to play the same note might sound ;)

"**Inv.**", short for Inverted, is a little abstract, note tuning is pushed more to the outside edges. This gives the ear a bit more of a sensation of 2 pitch centers (sharp and flat).

"**Random**" selects random detune amounts every time there is a new note-on. This might sound more natural in some situations, as of course humans don't play out of tune the same amount every note/time.

13.Wavetable Editor

The Wavetable editor is an advanced area of Serum for the purpose of creating and editing wavetables. Clicking the magnifier icon on the main panel waveforms will open the editor.

Enter the wavetable editor by clicking the magnifier icon in the oscillator A or B panel.

Thumbnail overviews

Thumbnails run along the very bottom of the Wavetable Editor. These allow you to perform a variety of functions:

Thumbnail overviews allow an easy way to view, select, and re-order the sub-tables that make up a wavetable.

• **Click** a thumbnail to view/edit a particular sub-table.

• **Click+Drag** left or right to re-order a sub-table to a new location. As you drag left or right past its neighboring sub-table, a yellow vertical cursor will appear, indicating where the table will end up being located, once you release the mouse button.

• S**hift-Click** a second table to select a range of tables (you'll see a blue selection). After this range selection is made, you can use functions such as "Remove Tables (multiselection)" and "Remove Tables (except selected)".

"Draw Tools" / Draw Shape Toolbar

The Tools are for drawing in the Big Wave display. They operate related to the grid sizes (set in the lower-right corner of the main waveform).

TOOL	NAME	DESCRIPTION
	Flat Line	The default and left-most tool. Think of this as a "plain draw" (it is, if grid is set to 0).
	Slope Up	Diagonal upwards line. Useful for drawing small saws to give your waveform a buzz.
	Slope Down	Diagonal downwards line. Useful for drawing small saws to give your waveform a buzz.
	Sine	Sine function. Useful for adding a pure harmonic at the grid size (e.g. if grid is set to 4, drawing with this will add an overtone two octaves above the fundamental).
	Half sine (valley)	Half-sine waveform.
	Half sine (peak)	Half-sine waveform.
	Curve up	Quarter-sine waveform.
	Curve down	Quarter-sine waveform.
	Interpolate linear	Rather than drawing like the prior tools, this will "connect the endpoints" in other words draw a straight line across the clicked grid segment. This can be useful for smoothing out the waveform (LPF).
	Interpolate curved	Similar to the above, however with a more gradual transition (not as much harmonics as linear / even more lowpass/smoothed).
	nudge	Move a portion of the waveform up/down. This can be used for clipping (if you click+drag vertically enough, the audio will begin to clip / flat-line.
	noise	Adds noise to the waveform. Dragging vertically up will add noise, while dragging down will reduce the noise.

FFT Area

The FFT area allows you to draw the waveform harmonics and their relative phases to one another,

The top row is the frequency bins. The left-most represents the fundamental or single sine wave at the oscillator's base frequency (in other words, the note you play). Each vertical bar to the right are the harmonics which make up the sound, (2:1 = octave, 3:1 = +19 semitones, 4:1 = 24 semitones, 5:1, etc). The second row is the phase offset for each of the harmonics.

You can ctrl-click (right-click Windows) for a context-menu. There is a separate menu for the Frequency and Phase bins:

MENU ITEM	DESCRIPTION
Clear All	Zeroes out all bins (silence)
Clear HF (bin X to end)	Zeroes all bins (silence) to the right of the the clicked mouse location. Since the high harmonics are being silenced, this can be thought of as a LPF.
Clear LF (start to bin X)	Zeroes all bins (silence) to the left of the the clicked mouse location. Since the low harmonics are being silenced, this should be thought of as a HPF.
Generate Saw	Generates all harmonics
Randomize Low X bins	Inserts random values for the left-most "X" number of bins. This will change the base tone of the sound while leaving the higher-frequency ("buzz") intact.
Randomize Low X bins (with half)	Same as above, but allows for a 50% amplitude for harmonics as well, this provides more tonal options.
Randomize All	Results in a noisy/buzzy tone. Useful as a starting point for further destruction with FFT drawing or other menu commands.
Create random series gaps	Randomly removes harmonics
Progressive Fade	Reduces the HF content gradually. Useful for bringing back a natural-sounding tonal balance to an overly bright or buzzy wavetable, without resorting to removing the highs altogether.
Shift Octave Up	Spreads all bins (1->2, 2->4, etc.) so that the resultant waveform is doubled in frequency.
Shift Octave Down	Same as above, but in reverse.
Draw Odd Harmonics Only	Prevents the mouse cursor from operating on the even-numbered bins.
Draw Even Harmonics Only	Prevents the mouse cursor from operating on the odd-numbered bins.

Formula Parser

This allows you to generate a wavetable (single sub-table, or entire wavetable set) from a math formula. A Formula presets menu exists as well, to get you started with some examples. See the formula parser chapter later in this manual for more information.

Menu Bar Commands

Above the waveform in the wavetable editor is a row of menus:

| Single | Process | Morph | Add/Remove | Sort | Import | Export |

Single Menu
process functions that will only alter the currently selected (visible) table.

"Single" menu Item	Description
Normalize	Adjust waveform to be maximum-scale amplitude
Remove DC Offset	Removes any DC (up/down bias to the waveform). Typically not needed.
Flip Vertical	Polarity invert (often somewhat mistakenly called "phase invert").
Flip Horizontal	Reverse the audio. Being a short looping wave, it probably won't sound 'reversed' in most situations.
Shift Horizontal to Zero Crossing	Nudges/wraps the wave data so the edges fall on a zero crossing. This isn't essential to do, but may be useful for applying fades or matching phase among multiple waveforms..
Init (silence)	Replaces the current table with silence.
Fade Edges (grid size)	Creates a fade in at the left edge, and a fade out at the right edge, based on the width of your horizontal grid size (set in lower-right)
X-fade Edges (grid size)	Similar to above, but instead of fading to the center line, the edges fade to each other.
Filter (grid size)	Low pass filter, completely removing upper harmonics. Lower grid size = more LPF. You can achieve the same LPF by using the FFT area ctrl/right click Clear HF Since this uses FFT it is infinitely steep, and you may notice some DC shift as all DC offset is removed.
Sample Redux at grid size	Provides a Sample rate reduction for a lofi sound. Grid size is not used in a literal way but as a means for specifying an amount (larger number = more SR Redux)

Process Menu

Process functions, as found in Single, but apply to all (1-256) sub-tables.

"Process" menu Item	Description
Normalize Each (gained separately)	Every table is individually normalized to its own peak maximum.
Normalize Same (max from all frames)	The entire wavetable set is scanned for a peak level, and the same gain amount is applied to each sub-table.
Remove DC Offset	Removes any DC (up/down bias to the waveform). Typically not needed.
Flip Vertical	Polarity invert (often somewhat mistakenly called "phase invert").
Flip Horizontal	Reverse the audio. Being a short looping wave, it probably won't sound 'reversed' in most situations.
Fade Edges (grid size)	Creates a fade in at the left edge, and a fade out at the right edge, based on the width of your horizontal grid size (set in lower-right)
X-fade Edges (grid size)	Similar to above, but instead of fading to the center line, the edges fade to each other.
Filter (grid size)	(See description previous page)
Remove Fundamental (HPF)	This removes the lowest frequency (pitch), effectively akin to zeroing the top-left bar of the FFT display (fundamental) but for all tables. The formula "z=(q>1)?in:0" would also yield the same result, or simply "(q>1)?in:0" if you want to only apply to the currently visible frame.
Sample Redux at grid size	(See description previous page)
Resize Tables to be Half (2x total)	Every table is divided in half, to become 2 tables. This doubles the number of tables so you should not use this feature if you have > 128 tables.
Resize Tables to be Double (½ total)	Every table is grouped with the following table. This is useful if you notice every 2nd table on a wav import looks similar (the frequency was detected an octave too high).
Create PWM from this table to all	Removes all tables except the current sub-table and creates a PWM shift across all 256 subtables. There is a realtime PWM effect under the main panel's Warp knob, but this way you can PWM and use another Warp effect.
Nudge all phases for fundamental to 50%	This is useful for preserving phase of the lowest frequency (fundamental) across various waveforms.

Morph Menu

Creates or removes interpolation sub-tables between the existing sub-tables. You must have more than 1 and less than 256 subtables.

After selecting a morph item from the menu, you'll notice that there are now 256 tables in the WTPos selector. This is because all empty wavetable slots are now filled with interpolated (blended) waveforms of the neighboring tables.

You'll also notice that the thumbnails no longer display "1, 2, 3" etc, but rather "1, 17, 33" (or similar). This is because the intermediate tables (e.g. 2-16 in this example) contain the newly-created morph tables. You'll see these intermediate tables in the waveform area as WT Pos is adjusted, but they don't exist as thumbnails (for your own comfort).

Once you exit the WT Editor back to the main panel, you'll also notice the interpolated tables will draw as grey in the 3d wave overview (whereas green is the 'real' subtables, and yellow is the current-selected sub-table).

"Morph" menu Item	**Description**
Morph - Crossfade	The interpolated tables are created by crossfading the neighbor tables together. This is the recommended default, and what traditional wavetable synths do.
Morph - Spectral	The spectral and phase content of the neighboring tables is used to re-synthesize the interpolated tables. This is what additive synthesizers do.
Morph – Spectral (zero fundamental phase)	Same as above, but the phase content of the fundamental is zeroed for all source tables. This way the lowest frequency will not shift/rotate between tables.
Morph – Spectral (zero all phases)	Same as above, but all phase content is discarded. This may alter the sound of the source content drastically, therefore sometimes undesirable, but will also make for the smoothest transitions between tables since no frequencies need to shift phase.
Remove Morph Tables	Get things back to how they were before interpolation was applied. "Undo" is possibly the better choice for undoing the morph however, in the case of the spectral (zero-phase) modes, since the two zero-phase choices (directly above) destructively alter the source tables.

Add/Remove Menu

Insert or Delete subtables. Init is also found here:

"Add/remove" menu Item	Description
Init All (Default)	All subtables/interpolation are cleared. The wavetable is brought to it's default state (single saw sub-table).
Insert at (current index)	Adds an additional sub-table to the wavetable set, inserted after the currently selected table. This is the same function as pressing the "+" portion of the [-][+] buttons below the waveform.
Remove (current index)	Removes the currently selected table frame. This is the same function as pressing the "-" portion of the [-][+] buttons below the waveform.
Remove: multiselection	Removes multi-selected table frames (shift-click in the thumbnail area to multiselect a range of tables).
Remove: beginning->selected	Removes the table frames from the beginning (1) through the selected table.
Remove: selected->end	Creates a fade in at the left edge, and a fade out at the right edge, based on the width of your horizontal grid size (set in lower-right)
Remove: All (except selected)	A "crop" feature, useful when you decide a single table or a certain range (multiselect) is all you want to keep.

Sort Menu

Re-order the existing tables automatically, based on a spectral property. Keep in mind you can 'manually' sort sub-tables by dragging the thumbnails left<->right at the bottom of the waveform editor window.

"Sort" menu Item	Description
Sort by spectrum (Peak Spect)	frames are sorted based on the which frame has the highest peak frequency bin.
Sort by spectrum (Average Spect)	Frames are sorted based on where the average spectral content exists (sum of all frequencies).
Sort by spectrum (Peak Amount)	Frames are sorted by the highest overall peak (concentrated frequency energy)
Sort by spectrum (Num w/ Spect)	Frames are sorted by how many frequency bins contain energy (spectral complexity)
Sort by spectrum (Highest w/ Spect)	Frames are sorted by the highest frequency bin to contain spectra. This should work well for filter sweeps, for instance.
Sort by spectrum (Fundamental amt.)	Frames are sorted by the amount of energy in the fundamental. This is a recommended sort to try first, when in doubt.
Reverse entire table order	The table order is reversed, useful if you decide you want the table order to become bright-to-dull instead of dull-to-bright, for instance.

Import Menu

Displays a dialog for selecting a file to load, which will overwrite the wavetable set you are currently editing. Keep in mind you can drag-and-drop to the waveforms, which is often a more ergonomic way to import sounds (for instance, you can typically use your host's file browser for finding a sound).

Please see the next chapter, titled "Importing Audio", for descriptions of the corresponding import choices.

Export Menu

Allows you to save the wavetable.

"Export All as Wav.." is identical to the Save "Floppy Disk" Icon on the main page of Serum. (.wav file, mono 32-bit 44100 Hz, 2048 samples per subtable. Additional header info for interpolation mode, interpolation tables are not saved as part of the wave data)

"Export Selection" saves a wavetable in Serum's format (.wav file with extra header information), just like the above export feature, except only the selected frame(s) are included in the resultant file.

"Export All 8-bit (.256)" is a file format used by some hardware modular

synths (Wiard, Piston Honda).

More export options may appear in the future if sufficient demand, but you should also be able to convert the .wav to your desired destination in a 3rd party application/utility.
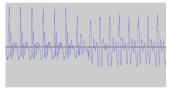
14.Importing Audio

 One method of getting good source material for wavetables is by importing audio data for the wavetables. There are several different methods for importing data, this is best thought of as a single-cycle or multi-cycle load.

Serum is **not a sampler** – at least in a traditional way. Don't expect to create ultra-realistic reproductions of acoustic instruments with Serum, it isn't designed with that in mind. With that said, in Serum, you can import from samples and manipulate the signal in ways that aren't possible in conventional samplers.

Multi-cycle waveforms – a brief overview

Most real sounds you come across (such as speech) is made out of multiple waveform cycles. In the human voice waveform pictured to the right, you can see this repetitive nature of the sound (which becomes the pitch we hear). If you have zoomed in on a waveform in your host, you have likely noticed this before.

waveform of human voice, approximately 15 cycles are pictured (this fact is noticed by counting the repeating peaks along the top).

Unlike a sampler which plays back through the audio as a single stream of data, when importing a sound into Serum, it attempts to slice the sound apart into the individual single-cycles. Because of the nature of audio, it tends to work best to select source sounds which are *monophonic* in nature, that is sounds which contain a single pitch (e.g. a single note is better than a chord). Because these single cycles become the basis of the oscillator, pitch information is effectively removed. In other words if you load a sound which has a pitch bend, the pitch bend will no longer exist.

Serum isn't attempting to be a sampler. Traditional sample playback, or even granular or additive synthesis do a better job of replicating a source sample in most situations. However wavetables have their own sets of strengths, which I would describe as a 'solid / fixed' sound which lends itself well to unison and wave manipulation (sync/fm/etc) without sounding flimsy.

With this in mind, Serum can do a fantastic job of importing the

waveform of other sounds. Fixed-pitch one-shot (monophonic) sounds, such as a one-shot sample of a synthesizer are the best choice for importing into serum, however speech and even complex sounds can yield some pretty interesting results if you're open to some experimenting.

Importing a Multi-Cycle wave file

To import a multi-cycle wav file, drag the audio file from your Finder/Explorer or Host Sequencer's *file browser* to the Waveform display (green waveform) on Serum's main window.

Note: Dragging files directly from the host arrangement windows or "region bins" is not possible in most hosts, but the host standard file-browser should work.   However many hosts have a way to show the (parent) sound file in the host file browser, and you should be able to drag from there.

These are the import options which appear when dragging a file to the waveform overview.

As you drag over the waveform, you will see a variety of choices. Where you 'drop' the file (release the mouse) determines the load method.

Serum will then analyze the sound and create a new wavetable set in RAM. The specifics of analysis depend on what method you chose for import (pictured above, also found in the "Import" menu of the WT Editor).

   The load 'choices' (pictured above) are as follows:

import: normal (dynamic pitch zero-snap)
This scans the audio file and builds a pitch-map. It then attempts to locate zero-crossings that fall near the pitch map. While this can work wonders on simple sounds, complex sounds unfortunately don't adhere to having sensible zero crossings, so you'll end up with some glitches at best. It is recommended to use this mode when you have a sound with a non-fixed fundamental (pitch bend or vibrato) **and** the sound is pretty simple, for example a sawtooth wave with little filter sweep/resonance.

import: normal (dynamic pitch follow)
Similar to the above, this also will build a pitch map and import a varying-sized segment of audio for each sub-table, based on the analyzed pitch. Unlike the above, pitch-follow import does **not** attempt to locate zero crossings. This means it is better suited for complex sounds, such as a source sample that might have a touch of

chorus/unison, resonance, or background noise/notes.

Import: constant frame size (pitch avg)
*If in doubt... try this one first!

This is typically the best choice when a sound has a fixed frequency, such as a one-shot from a synthesizer (in other words you hear it as a perfect or near-perfect constant-pitch, with no pitch bend or vibrato to speak of).   In this mode, Serum will analyze the entire file for an average pitch, and then uses this number of samples as the import length. Because some sounds contain half-cycles, silence, multiple notes, among other artifacts, serum might not guess the desired pitch. Fortunately, Serum will display the number of samples it is using per-frame in the WT Editor Formula area (read on...) and will switch over to this 'fixed' value found from analysis unless changed or cleared from the formula field.

Advanced Import: constant frame size (manual entry)

Due to the complex nature of audio signals, using a pitch average may not always be perfect. Sometimes you may wish to specify an exact number of samples for each cycle (which you can measure yourself in a sample editor).

In the WT Editor, there is a formula parser (which is described in the next chapter), a text box below the large waveform which reads "(enter formula)". Although this is text field's primary function is for creating waveforms from functions (see next chapter for that), you can also type either:

• a 1-4 digit number in this box (e.g. 1024) to instruct Serum to split the dragged sound file into segments at this number of samples.

• a MIDI "note name" (e.g. B0, C#2, D6, etc) and Serum will convert to the appropriate **nearest*** number of samples for you for the split (rounds to the nearest whole number of samples). This MIDI note/sample conversion assumes a 44100 Hz source sound file.   If your source sound file is not 44100, the number of samples is still held true, but the displayed frequency/note information on the formula text area should be disregarded.

*since musical note frequencies are not often perfectly divisible by a sample-rate, many pitches will have some visible 'drift' on import. See the section named "For Best Results" later in this chapter.

If you wish to clear this value, so Serum will once again present the choices as it does by default, simply type a zero ("0" without the quotes) in the Formula text field. Once you type the 0 and press

Entering zero (0) in the WT Editor so Serum no longer divides a file by a pre-determined length on import.

Enter, Serum will once again display "(enter formula)" and dragging a file in to the waveform will once again display the options mentioned at the start of this chapter.

Protip: Since serum is not playing audio back quite like a sampler ("looping tiny portions of the sample" is a simple way to think of it) you might here undesirable buzz from subharmonics or other factors causing the beginning and end of the wave cycle to not line up completely. Therefore after importing with one of the above methods, try using a fade command from the Process menu (fade 16 samples, or fade grid size) to give you a less buzzy-sounding result!

FFT 256 / FFT 512 / FFT 1024 / FFT 2048

Unlike the other import modes, which import + divide (and possibly stretch) chunks of the original waveform, the FFT modes are a spectral import. In other words, these import modes divide the source audio into small snippets of time, and analyze the spectral content. One way of thinking of this is a "blurred averaging of the frequency content".   This can be very useful for sounds like drumloops, speech, and other material where you want the flavor of the sound for abstract purposes. Serum isn't a sampler, and these import modes aren't going to leave you with an exact replica of your source audio material.   The larger the number, the longer the time period for each slice of source material. Therefore a larger number means less **time** resolution, however it also means a higher **frequency** resolution.

[ADVANCED] Text file override for getting specific results

[programmer's note: This is not a very common feature to need, or user-friendly, but since it exists, I figured I would document it, rather than keep it a secret!]

When dragging an audio file for import, Serum looks to the file location of the file you dragged. Serum then looks for a file with the same *name* but with a .txt file *extension* instead (MySound.wav   Serum looks for MySound.txt). If it does not see this sound file, it also looks for FolderInfo.txt in the folder. (FolderInfo.txt is the same idea, but will apply to any sounds in the folder, this prevents you from having to make the same text file for every wav file in the folder).

The anatomy of the text file is pretty simple. On the first line is a number contained in brackets, such as:

[512]

This specifies the number of samples-per-cycle of the audio file. This can be a number between 32 and 9999.

The second line (optional) will specify whether or not intermediate interpolated tables should get created. The default is yes, if you wish there not to be you just need an [n], e.g. [no interpolation].

Here are a couple example text files, first for a 786 samples-per-cycle waveform with crossfade interpolation:

[768]

[interp]

(the second line is optional as it is the default)

Here is an example of 1024 samples-per-cycle with no interpolation:

[1024]

[no interp]

For best results...
Serum internally uses 2048 samples for each single-cycle. This makes 2048 the ideal number of samples to use as a wave cycle, if you plan on generating sounds from synthesizers, vocoders, etc. to import in to Serum, because no resampling of the source audio will be needed. This 2048 samples-per-cycle works out to 46.875 Hz at 96000 Khz sampling rate (96000/2048 = 46.875), which is F#0 +24 cents (note the octave may differ in various hosts and synthesizers, your best bet is to render and measure in a program that allows you to select in samples, such as the freeware sound file editor Audacity). Type 2048 in the Formula field of Serum's wavetable editor prior to drag-import and Serum won't have to alter the sample data at all.

Single-Cycle waveforms
To load a *single-cycle* wave, (or load a short sound to be interpreted as such), drag an audio file (from Finder/Explorer or your Host Sequencer's file browser) to the desired Thumbnail overview in the Sample Editor window. This will just replace the waveform in that particular table, and you'll still have the rest of the wavetable set in the other table slots.

To load multiple *single-cycle* wavetables at once, you can drag them at once to the waveform overview on Serum's main display window. Because Serum will notice that you dragged multiple files at once, it will then treat them like single-cycle waveforms. This will replace your current wavetable with a new one comprised of these files you just dragged.

Import sub-menu in WT Editor

The same import options you get from drag-and-drop also exist in the WT Editor sub-menu titled "import". This way if you would rather get a standard file-browser window instead of drag-and-drop, you can use the items in the menu.

[ADVANCED] Import Tutorial: Creating a sound specifically to import in to Serum

Perhaps we want to import a waveform/sound from another soft synth into Serum to use as a wavetable, and would like maximum quality.

When you have the luxury of choice, it is best to use low-frequency notes. the ideal is 2048 samples-per-cycle, because that is the internal storage size and means no special interpolation is used/needed to get the sound to this size. Alternatively (and slightly easier to set up as no fine tuning of pitch is needed) low C at 44100 works well (33 Hz, 1349 samples per cycle).

In this example we will export a 2048 samples-per-cycle wave from another software synthesizer and import it into Serum.

1) Determine the appropriate source note (pitch)

Open Serum's WT Editor, and type 2048 into the formula text box. Serum replies with:

split at: 2048 samples (22 Hz, note: F-1 and -22 cents)

Now we know we should play a low F with -22 cents will give us our desired length of 2048 samples-per-cycle.

2) Render a file playing this pitch from the softsynth

2A) Create the F-1 note in your host sequencer's piano roll on the softsynth track (octave naming may be different as there is unfortunately no standard, so you should probably use a spectrum analyzer to verify ~22 Hz is the fundamental, not ~11 Hz or ~44 Hz). Duration of note: Since Serum can handle 256 wave cycles, we can render this 22 Hz file for up to 11.6 seconds, which is about 8 bars at 165 BPM. If the sound has very little modulation/change to your source sound, there is no need for such a long note - you can probably determine by ear when the source sound has stopped changing / no need to have a longer note.

2B) set the fine tuning to -22 cents. How you do this varies in different softsynths, usually it is set per-oscillator, but some synths have a global tuning. If you can't do that in the source synth, perhaps you can add the appropriate amount of "pitch bend" down as a MIDI

message (22% down if pitch bend range is 1 on the source softsynth).
A=440   approximately A=434.5 Hz is another potential option.

      2C) Render/export this note from the softsynth in your host as audio (similarly, you can "Freeze" + "Flatten" to get this softsynth note as an audio file).

3) Drag this render in to Serum

Locate the audio file and drag in to Serum. If the "2048" is no longer visible in Serum's Formula display (You closed the Serum window since step 1), you should re-type that in. When you drag the waveform over Serum's waveform, it should display the same text we see in the WT Formula ("*split at: 2048 samples (22 Hz, note: F-1 and -22 cents)*").
When you release the mouse button, the audio is imported. The source audio file (render) is no longer needed at this point, unless you wish to re-import it into Serum with different settings.

4) Clean-up and save

• You will likely have blank or unnecessary tables at the end of the wavetable (click on thumbnails and listen/find the first WT you don't want, and select "Remove: Selected->End" from the Add/Remove menu in the WT Editor.

• You might want to apply fades to the tables (use one of the fade or x fade options from the WT Editor's Process menu) to eliminate any buzzing artifacts from frame edges. Ideally this isn't needed, and if your source sound has no subharmonic content, it may not be needed.

• You can listen to the entire WT set by playing a note and dragging the WT Pos number box in the WT Editor, or similarly the WT Pos knob on the main panel for the oscillator. For an automatic playback, you can alternatively use a mod source for this, e.g. an envelope with long attack, or an LFO with upward saw shape, to modulate the WT position.

15.Formula Parser

Type some math, make a sound.

The Formula Parser (a type-able text field located in the Waveform Editor) allows you to enter math functions in order to plot tables generatively.

This is a fairly advanced feature, and something most people probably have no desire to explore. However if you want to generate waveforms from thin air, then perhaps now is the time you'll finally wish you paid attention in math class!

Fortunately there are many formula presets included, and it is pretty easy to paste in formulas provided by other users (I'm sure a nice list will get made in the Serum registered user forum!)

Basic Functions

Name	Arguments	Explanation
sin	1	sine function
cos	1	cosine function
tan	1	tangens function
asin	1	arcus sine function
acos	1	arcus cosine function
atan	1	arcus tangens function
sinh	1	hyperbolic sine function
cosh	1	hyperbolic cosine
tanh	1	hyperbolic tangens function
asinh	1	hyperbolic arcus sine function
acosh	1	hyperbolic arcus tangens function

Name	Arguments	Explanation
atanh	1	hyperbolic arcur tangens function
log2	1	logarithm to the base 2
log10	1	logarithm to the base 10
log	1	logarithm to the base 10
ln	1	logarithm to base e (2.71828...)
exp	1	e raised to the power of x
sqrt	1	square root of a value
sign	1	sign function -1 if x<0; 1 if x>0
rint	1	round to nearest integer
abs	1	absolute value
min	var.	min of all arguments
max	var.	max of all arguments
sum	var.	sum of all arguments
avg	var.	mean value of all arguments

Built-in binary operators

The following table lists the default binary operators supported by the parser.

Operator	Meaning	Priority
&&	logical and	1
||	logical or	2
<=	less or equal	4
>=	greater or equal	4
!=	not equal	4
==	equal	4

Operator	Meaning	Priority
>	greater than	4
<	less than	4
+	addition	5
−	subtraction	5
*	multiplication	6
/	division	6
^	raise x to the power of y	7

Constants and Variables:

sel – refers to the current ("old") waveform value of the selected table (does not change to each table when using a multi-table formula, in other words the formula contains y or z) just refers to the selected table when processing the formula begins).

rand – a random number from -1.0 to 1.0, stays the same for all tables

Name	Explanation
pi	pi (3.14159265897932384626 4338)
e	e (2.71828182818281828 1828)
w	current time-value getting plotted, from 0.0 to 1.0 same as (x+1)/2
x	current time-value getting plotted, from -1.0 to 1.0
y	current table number, from 0.0 to 1.0 ***See Below**
z	current table number, from -1.0 to 1.0 ***See Below** equivalent: (y*2)-1
q	When a 'q' is preset in the formula, the function plots to the FFT bins instead of plotting the to the waveform display. q iterates from 1 to 512 for the respective harmonics/bins. **
in	the current ("old") visible waveform value of the plotting table (changes to each old table, if using 'y' or 'z' which plots all

Name	Explanation
	tables)
`sel`	Similar to "in", but only the currently selected wavetable (does not change to each table, uses the selected table when processing the formula begins).
`rand`	a random number from -1.0 to 1.0, stays the same for all tables (precalculated for every time position).

* 'y' and 'z' refer to the current table number. Thus, when y or z is in the expression, *all tables* for the current oscillator will be regenerated!

** 'q', as mentioned above, renders the function to FFT instead of to the waveform display. There is no reason to include 'x or y 'in a formula containing 'q'

Formula Presets Menu
There is a formula menu which can get you started with examples. This is a good way to learn the capabilities, as well as make your own presets. As soon as you select a formula, it is entered into the formula text-area automatically and the waveform is calculated.

Single-cycle formulas submenu ("Singles")

Formulas in the "Singles" sub-menu do not contain 'y' or 'z' variables, so when one of these is selected, only the currently selected table will be replaced/generated with a formula.

Multi-cycle formulas submenu ("Multis")

Formulas listed in the "Multis" menu do contain 'y' or 'z' variables, so when selected, the entire wavetable will be replaced.

Saving your own formulas

There are two files in SerumPresets/System/

FormulaUserMultis.txt and FormulaUserSingles.txt – these files, if not empty, will add additional submenus "User Singles" and "User Multis" to the Formulas list. Serum plug-in must be de-instantiated and re-inserted for formula menus to refresh. The text files have a format as follows:

[formula1][formula1 name]
[formula2][formula2 name]

Refer to the files FormulaFactorySingles.txt or FormulaFactorySingles.txt for examples.

Note: While you can technically also edit FormulaFactorySingles.txt and FormulaFactoryMultis.txt, it is recommended to use the User Files instead. This way you won't accidentally overwrite your personal formulas when installing an update, for instance.

16.[appendix] File Structure on disk

Here is a summary of the data saved to disk from Serum.

"Serum Presets" Folder

The "Serum Presets" folder contains everything Serum reads and writes, except the preferences file. This folder installs by default to:

(OSX) /Library/Audio/Presets/Xfer Records/

(WIN) /Documents/Xfer/

This folder contains the following subfolders:

• Presets
In this folder are sub-folders containing the presets (.fxp files) that you see in the main menu (factory presets, and any ones saved in User). You can create your own subfolders in Presets, but Serum does not scan deeper (no sub-sub-folders will be scanned).

• Tables
In this folder are sub-folders containing the Tables (special serum-saved .wav files) that you see in the main menu (factory presets, and any ones saved in User).   You can create your own subfolders in Presets, but Serum does not scan deeper (no sub-sub-folders will be scanned).

• LFO Shapes
This folder contains the .shp (Shape) files which are the same file format as the shape files used in the LFOTool plug-in. The LFO Shapes will appear in:

– the LFO section of Serum (lower-center, Folder icon)

– the waveshaper (FX Section, and X-Shaper effect type is selected and waveshaper graph is displayed).

– The Remap editor (Warp mode is set to Remap and the graph is displayed).

• Noises
This folder contains subfolders for the Noises which appear in the noise menu. You can add your own noises to these folders but it is not

recommended / do sparingly and at your own risk: Mostly because, unlike wavetable data, Noise data is not stored in the preset so you will need to keep your custom noises if you wish to open presets/songs in the future. Otherwise you will be notified of the missing noise sound.

• System

The System folder contains the four Formula files (FormulaFactoryMultis, FormulaFactorySingles, FormulaUserMultis,FormulaUserSingles) . You are free to edit any of these 4 text files, but the intent/recommendation is to edit the User files and leave the Factory ones alone. See the previous chapter on Formulas for more information.

Also in the System folder is a MIDIccMaps folder, which is where any MIDI cc maps you create get stored.

Lastly is a User.dat file which holds your registration information.

Serum Preferences file

The other file on the hard disk is the Serum Preferences file named Serum.cfg - this file exists in a fixed location:

(OSX) ~/Library/Preferences/Serum.cfg

(WIN) C:/Users/You/AppData/Roaming/Xfer/Serum

This file will automatically get re-created by Serum if it goes missing, so you can reset your preferences to the factory defaults simply by deleting this file.

This file holds the preferences accessible in the Preferences page, the last known path to the Serum Presets folder, and sometimes it may hold power user options. At the time of this manual, the only reason to edit this file would be if you want to add a default user name to the Init preset. The area for this would be up top, between the curly braces { }. So you can edit this Serum.cfg in a text editor, type your name between the curly braces, and now any presets made from scratch will have your name listed as Author.

17.[appendix article] Making wavetables

What makes a good wavetable?

Like a good song it's about getting that careful balance between consistency *and* contrast. Too much consistency/repetition and things won't feel like much of a journey (boring/constant). Too much contrast and it is difficult to draw connection between random events (e.g. turning the dial on the radio feels disconnected and hard to comprehend as intentional).

The two Categories of wavetable types

For good tables, I separate them mentally in to 2 basic categories:

A) many tables, but correlated This would be typically a sound which has a lot of frames to it, but they all fluidly work together- think a pluck string which decays across many cycles. There might be a lot of individual cycles, but they all go together well / relate to each other a lot. Another example would be a synthesizer sample of oscillator sync. You would want many cycles so it doesn't feel too 'steppy', however all of the tables feel like part of a collection (a similarity / sweep across them).

Or else....

B) few tables. Most of the factory tables in popular wavetable synths are only made up of 4-5 cycles. This allows for some variety within the table, but not too many waveforms that it feels disconnected/crazy.

Table ordering

It typically makes sense to have the tables progress from dull to bright, or vice-versa. You may have a situation where you want 'peak' spectrum in the middle somewhere, but probably not. You can drag the thumbnails

at the bottom of the waveform editor to re-arrange them. This way when you move the WT Pos, it feels like you are traveling in a straight line, instead of some zig-zag fashion, spectrally speaking.

Interpolation

If we have 4 tables and automate (e.g. move knob) on WT Pos, we'll hear 4 discrete tones and it will immediately jump from one to the next. This is typically undesirable. In the Editor you can select Morph->Table Morph (Crossfade). This is what many wavetable synths do automatically, in every preset wavetable. However there are times when you may prefer the discrete (non-interpolated) waveforms.

Making tables from 'scratch'

There are several different ways to 'make' source audio for the WaveTable:
     • draw. Grid size in lower-right determines snapping. Try different sizes such as 6 or 12 to bring out +7 and +19 harmonics.
     • draw in FFT bins at top. ctrl-click in bins for a popup menu with more options (random etc).
     • move drawing up to FFT bins (Curvy arrow on top left) and then add a new table (> button in lower left). Then make some adjustments and repeat. Again, about 4 tables tends to be a popular number of frames in many other software synthesizers.

Filter Types and Var parameter function

Filter Type	Description	Var Function
Mg 6/12/18/24	Ladder-Style Lowpass filter	"Fat" amount
Low 6/12/18/24	State-Variable Lowpass Filter (SVF)	"Fat" amount
High 6/12/18/24	State-Variable Hipass Filter (SVF)	"Fat" amount
Band/Peak/Notch		"Fat" amount
LH/LB/LP/LN/HB/ HP/HN/BP/BN/ PP/PN/NN	Dual SVF filters: first letter is primary, 2nd letter is secondary (e.g. BP = Band+Peak). Reso is linked, second cutoff is the Var param.	Cutoff Frequency "2" (for the 2nd of the two filters)
LBH/LPH/LNH/BPN	Morphing SVF filters (e.g. Lowpass<->Bandpass<->Hipass)	Morph (between the three filter states)
CombL /FlangeL / PhaseL	Comb/Flanger/Phaser with a Lowpass filter in the internal feedback circuit. [***]	LP Cutoff
CombH /FlangeH / PhaseH	Comb/Flanger/Phaser with a Hipass filter in the internal feedback circuit.[***]	HP Cutoff (feedback filter)
CombHL / Flange HL / Phase HL	Comb/Flanger/Phaser with a Hipass filter in the internal feedback circuit. [***]	HL Width (band separation)
EQ	Shelf (L / H) / Peak EQ	dB Gain
Combs/ Allpasses/ Reverb		Damping
French LP	A unique distorting Lowpass Filter.	Boeuf
German LP	A clean "Zero-Delay Feedback" Lowpass filter.	(none)
Add Bass	A phase-rotated lowpass filter with a touch of drive. Not a typical synth filter, but maybe you'll find a use!	Thru (phase shifted dry)
Formant 1-3	Formant 'vowel' filters. Cutoff knob morphs between the formants	Formant shift

[***] = for best result for Flanges menu items, set mix knob at 50%

This page unintentionally left blank. Oops.
Perhaps you can draw a picture below and use it as album artwork!