Automaton User's Guide

Audio Damage, Inc.

Release 1.1

28 December 2011

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

© 2008, 2011 Audio Damage, Inc. All rights reserved.

Credits

Software Design and Construction, Documentation
- Chris Randall
- Adam Schabtach

Field Testing
- Wade Alin
- Dean Dunakin
- Steve Hamann
- Jeff Laity

Made Possible By
- Elle
- Tracie

Fuzzy Logic
- Alex
- Chica
- Pablo
- Widget

Audio Unit version implemented using Symbiosis from NuEdge Development.

Audio Damage would like to thank Mirek Wójtowicz, whose website at http://www.mirekw.com/ca/index.html provided valuable information and inspiration during the development of Automaton.

License Agreement

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

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

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

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

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

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

This warranty does not apply if the software (a) has been altered, except by Audio Damage, (b) has not been installed, operated, repaired, or maintained in accordance with instructions supplied by Audio Damage, (c) has

been subjected to abnormal physical or electrical stress, misuse, negligence, or accident, or (d) is used in ultrahazardous activities.

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

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

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

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

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

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

System Requirements

To use Automaton, you'll need a Steinberg VST-compatible host application which conforms to the VST 2.0 specifications, and a computer capable of running it. For the Audio Unit version of Automaton, you'll need an application capable of hosting Audio Unit plug-ins, and an Apple computer with an Intel CPU capable of running it.

We support the use of Automaton under Microsoft Windows XP, Vista, or Windows 7; and Apple OS X version 10.5 or newer.

Installation

Double-click the Automaton Installer icon and follow the instructions. During the installation process the installer will ask you to enter your registration code. Your registration code uniquely identifies your purchase, and you will need it if you need to reinstall your plug-in (for example, after upgrading to a new computer). Keep a copy of the code in a safe location and please don't share it with your friends. We're delighted if you like our products so much that you want to share them, but please ask your friends to buy their own copy so that we can keep making new products.

On OS X, the installer installs both the AudioUnit and VST versions of Phosphor by default. If you click the Customize button you can choose to not install one of these formats. If you're running a 64-bit version of Windows, you can choose to install either a 32- or 64-bit version of Phosphor, or both. On OS X, both the AudioUnit and VST versions are combination 32/64-bit plug-ins.

To un-install from OS X, simply delete the plug-in from your VST folder, which is usually located at `/Library/Audio/Plug-Ins/VST/`, and your Audio Units folder, which is located at `/Library/Audio/Plug-Ins/Components/`. To un-install from Windows, delete the file named `Automaton.dll` from your VST folder(s).

Online Videos of Automaton

We've created several videos showing Automaton in action. You might find them to be a useful addition to this manual while learning about Automaton. The videos can be found at http://vimeo.com/audiodamage

Introduction

The heart of the Automaton plug-in is a mathematical construction called a *cellular automaton* (pluralized *automata*). The study of cellular automata is a wide field with applications in computer science, mathematics, biology, and physics. We won't attempt to provide anything more than a very brief introduction here. If you're interested in learning more about cellular automata, the Wikipedia entry is an excellent place to start: http://en.wikipedia.org/wiki/Cellular_automata

A cellular automaton consists of a rectangular grid of *cells* which we'll call the *world grid* or simply the world. Each cell can be in one of two *states*, usually described as "alive" and "dead". Cells change from one state to the other when a set of *rules* is applied to all of the cells simultaneously. Typically (and in the case of Automaton) the rules work by counting the number of live neighbors that a cell has, and setting the state of the cell based upon some function of the number of its neighbors and its current state. Only the cells immediately adjacent to a cell are considered neighbors, so each cell has eight neighbors. The same set of rules is used for each cell.

When the world is updated with a new population of cells that has been calculated by applying the rules to every cell, we say that a new *generation* has been created, and that the population of the world grid has evolved. Automaton's generations occur in sync with your plug-in host's transport, e.g. a new generation might be calculated every eighth note.

Although the rules for a cellular automaton are usually quite simple, the manner in which the world's population evolves over a number of generations can be surprisingly complex and unpredictable. Geometric patterns form and dissolve, order turns to chaos and back to order, stable groups of cells remain unchanged and other groups pulsate and move around while spawning new groups, and so on. Automaton harnesses this mixture of chaos and predictability to provide an inspirational new signal processor—a tiny world of effects processing evolving inside your DAW.

Operation

Automaton is a complex plug-in so we'll describe each of its features and controls in detail, but first we'll provide a quick overview and introduce some terminology. Here is a screen shot of Automaton in action:

The rectangular grid at the top of Automaton's window is the cellular automata world grid. Each square represents one cell. If the square is filled with white, the cell is alive. The population of live cells changes each generation when Automaton uses one of several rule sets to calculate how the world evolves.

The cells in the world do not directly affect the audio flowing through Automaton. Automaton has four signal processors, named Stutter, Modulate, Bitcrush, and Replicate. These processors operate independently and are arranged in series. The row of colored buttons below the world grid lets you choose which processor's controls (knobs, sliders, and switches) are displayed in the panel near the bottom of the window.

The colored dots in the world grid are triggers. The color of the triggers corresponds to the colors of the processor selector switches. A signal processor is turned on when there is a live cell in the same square as one of its triggers. We refer to a trigger that is in the same square as a live cell as an active trigger. If there are no active triggers for a processor in the world, you won't hear that processor; and if there are no triggers at all, Automaton doesn't change the audio flowing through it.

For example, if the image on the right showed the only contents of the world grid, you'd hear the Modulate and Replicate processors but not the Stutter or Bitcrush processors.

Automaton can be used in either a stereo or mono context in your host DAW software. If you use a stereo input, the left and right channel separation of the input signal is preserved in the output signal. Automaton can be used either as an insert effect or a send/return effect and has several different mixing operations to enhance its use in either context.

Automaton's knobs respond to vertical mouse-pointer movements; that is, click on a knob and drag upwards and downwards to rotate the knob. If you're using the VST version of Automaton, your host software may override this behavior.

Working in the World Grid

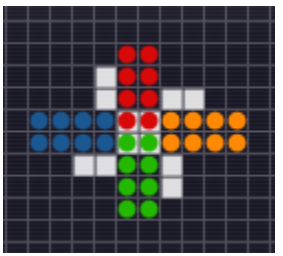

To add live cells to the world, click the mouse pointer on the desired square. You can click and drag the mouse around to add many cells at once. To remove cells, click on them with the mouse. As you would probably expect, you can remove many cells by clicking on a live cell and then dragging the mouse pointer around. Note that live cells obscure, but do not remove, triggers. You can also remove all of the live cells from the grid by clicking the **CLEAR** button in the Sequencer panel. (We'll discuss Automaton's panels in more detail shortly.)

To add triggers to the world grid, click the mouse pointer on one of the four processor selector switches, then click and drag in the grid. This gesture adds triggers for that processor to the grid, as you can see from the color of the triggers and the current processor selector switch. You can remove triggers in the same way as you remove live cells: click the mouse pointer on an existing trigger and drag the pointer around. You can also remove all of the triggers for the current processor by clicking the **CLEAR** button on the right side of the processor's panel. There can be only one trigger in a square in the grid, but there can be any number of triggers for each processor.

Note that when the Sequence panel is active, live cells in the world grid cover any triggers underneath them. If any of the effects processor panels are active, all of the triggers are drawn on top of live cells.

Automaton’s Control Panels

Next we’ll look at the panels of controls for each of Automaton’s signal processors, and for the controls which affect the operation of the world grid and how Automaton mixes its output with the incoming signal. Click on the corresponding button near the center of Automaton’s window to reveal each panel.

Common Controls

Many of the knobs in the signal processor panels have two adjacent sliders. These sliders modify the values set by the knobs. The first slider, labeled **?**, randomly changes the value set by the knob. The further upwards you move this slider the more the associated parameter will vary randomly from the value set by the knob. In most cases the randomization slider will increase the parameter's value, so you may notice that it produces a greater range of variation if the knob is set to a low value.

The second slider, labeled **P**, is the Population Modulation slider. This slider varies the value set by the knob by counting the number of active triggers for the processor. The amount that the Population Modulation slider affects the parameter depends upon several things: the total number of triggers placed in the world grid for the processor, the number of active triggers in the current generation, and the height of the slider itself. The more active triggers there are, the greater the effect of the Population Modulation. The effect is also scaled by the position of the slider: the higher the slider, the greater the effect. Population Modulation always increases the value of the parameter above the value set by the knob, so again you'll usually hear the influence of this slider more clearly if the knob is set to a relatively low value.

That all sounds complicated but it boils down to this: the greater the number of active triggers, and the higher the setting of the **P** slider, the more that the associated parameter will be increased above the value set by the knob. Thus the **P** slider allows the activity in the world grid to not only activate Automaton's effects processors, but also change the way they sound.

On the right-hand side of each processor's panel you'll find three controls labeled **CLEAR**, **SOLO**, and **MUTE**. Click the **CLEAR** button to remove all of the processor's triggers from the world grid. (If you want to remove all of the triggers for each processor and all of the live cells, click the **CLEAR ALL** button in the Sequencer panel.)

The **SOLO** and **MUTE** switches act in a manner similar to that of the solo and mute controls commonly found in mixers. When a processor's **SOLO** switch is illuminated, you will hear only that processor in action; the other processors will be silenced. Only one processor can be soloed at a time. If you turn on the **SOLO** switch for, say, the Stutter processor, switch to the Modulate processor and turn on its **SOLO** switch, you will then hear only the Modulate processor. The **MUTE** switch silences the processor. More than one processor can be muted at a time.

The **SOLO** and **MUTE** switches toggle on and off. Click them to switch them between their on and off states.

Sequence

The Sequence panel contains all of the controls that affect how Automaton's CA world evolves, including how its generations are synced to your host's transport and your music's tempo.

The ***SEQUENCE RESET*** knob determines how long the current world population evolves before being reset to its initial state. The range of this knob is a quarter of a measure up to four measures, or an infinite number of measures. For example, if you set this knob to two measures, the world will evolve for two measures, return to its initial state, evolve again in the same way for two measures, over and over again. If you set the ***REPEATS*** knob to its maximum setting, denoted by "INF" for infinite, the world will evolve without ever being reset to its initial state.

So, what do we mean by the *initial state*? The initial state is determined by either of two actions. First, when you load a preset into the plug-in (using whatever method is normal for loading presets in your host software), the state of the world set by the preset becomes the initial state. Second, when you turn cells on or off by clicking and dragging in the world grid, the state of the world when you release the mouse button becomes the initial state. This means that it's much easier to tell what the initial state of the world will be if you stop your host's transport while editing cells. If you don't stop the transport, the world will continue to merrily evolve as you drag the mouse around and the initial state will be whatever the world happens to

contain when you release the mouse button. In other words, establishing the initial state while the transport is running is something of an exercise in hitting a moving target.

The ***SEQUENCE LIFESPAN*** knob controls how rapidly the world evolves. That is, it controls how often a new generation of cells is calculated within the world grid. The generations are synchronized to your host's transport, so the Lifespan knob works in units of beats. It has a range of a 1/32 note up to a quarter note, and dotted and triplet values are available.

The ***SEQUENCE RESET*** and ***SEQUENCE LIFESPAN*** knobs have related roles. The ***RESET*** knob controls the duration over which the world will evolve before being reset to its initial state, and the ***LIFESPAN*** knob controls how rapidly the world will evolve during that duration.

The ***RULE SET*** popup menu lets you choose between several different rules which govern how the world evolves from one generation to the next. The rules are as follows:

Life

In the Life rules, if a living cell has two or three live neighbors, it stays alive in the next generation, otherwise it dies. If an empty cell has exactly three live neighbors, a new cell is born in the next generation. CA worlds that use the Life rule set tend to become less populous or die out althogether as the world evolves, but there are some cell patterns that will grow indefinitely or repeat themselves periodically. The Life rule set uses the set of rules invented by John Conway. Life is easily the most well-known cellular automaton rule set and was introduced to the (real) world in 1970 via an article in *Scientific American* magazine written by Martin Gardener[1].

Sequencer

The Sequencer rule set is a specialized rule set that makes Automaton's world grid act like a familiar step sequencer. Live cells move from left to right, advancing one square with each new generation. For example, you can set the ***SEQUENCE RESET*** knob to one measure, the ***LIFESPAN*** knob to 1/16, and draw a vertical strip of live cells along the left edge of the world. When your host's transport is running, the vertical strip of cells will move one space to the right with each sixteenth note and jump back to the left edge at the beginning of each measure. Then you can place triggers on the world grid

[1] Yes, this effectively makes it ancient history as far as computers are concerned. The article suggested using checkers and a checkerboard to iterate generations by hand. No, the author wasn't kidding.

to activate Automaton's effects at sixteenth-note intervals as you desire. The ***SEQUENCE WRAPPING*** switch (described below) has no effect when this rule set is chosen.

Replicator

Worlds evolved with the Replicator rule set have the interesting property that any pattern of cells is eventually replaced with several copies of the pattern (provided that the world is big enough). In this rule set, any cell square will contain a live cell in the next generation if it has an odd number of live neighbors (i.e., one, three, five, or seven live neighbors). With the Replicator rule set active, Automaton's worlds tend to become denser as they evolve, sometimes jumping back to relatively sparse populations at unpredictable intervals. The Replicator rule set is attributed to Edward Fredkin, one of the early CA researchers and also creator of the somewhat legendary Triadex Muse algorithmic music-generating device built during the 1970s.

Serviettes

In the Serviettes rule set, every live cell dies in the next generation. New cells are born in empty squares if the square has two, three, or four live neighbors. Cell patterns governed by the Serviettes rule set tend to grow into elaborate, symmetric patterns that sometimes resemble lace (hence the name).

Gnarl

Of the rule sets available in Automaton, the Gnarl rule set is the simplest to describe: A square will contain a live cell in the next generation, regardless of whether it is currently live or dead, if it has exactly one live neighbor. Despite its simplicity, Gnarl can generate cell populations that rise and fall in density in unexpected ways. This rule is attributed to Kellie Evans.

The ***SEQUENCE WRAPPING*** switch affects what happens to cells at the edges of Automaton's world grid. If the ***WRAPPING*** switch is off, the edges are boundaries, as though the world grid was surrounded by a vast, lifeless, digital wasteland. Hence, when applying a rule set to calculate the next generation, a cell in any of the four corner of the world grid can have at most three live neighbors, and one on any of the four edges can have at most five live neighbors.

If the ***WRAPPING*** switch is on, the world grid is treated as though its left edge meets its right edge, and its top edge meets its bottom edge. To envision how this might be done, imagine drawing a grid on a donut with a felt-tip pen[2]. Every square has eight neighboring squares, regardless of its location on the grid.

With ***WRAPPING*** turned on, some interesting things can happen if the cell population is allowed to evolve indefinitely. For example, in a world dictated by the Life rule set, there are patterns of cells known as gliders which move across the grid as they cycle through a series of states. If the ***WRAPPING*** switch is off, gliders either disappear or turn into a motionless block of cells when they hit one of the edges of the world. But if wrapping is turned on, gliders can sail along indefinitely, moving off one edge of the world and reappearing on the opposite edge.

The buttons labeled ***CONTROL*** in the Sequence pane do several different things. The ***CLEAR*** button, when clicked, erases all cells in the world grid. The ***CLEAR ALL*** button erases all cells and all triggers en masse.

The ***RAND RST*** switch, whose name is short for ***RANDOM RESET***, when turned on, causes the world to be erased and filled with randomly placed live cells when the sequencer resets (i.e. at the interval set with the ***SEQUENCER RESET*** knob). If you turn this switch on, the world will be replaced with a different configuration of cells every time the sequencer restarts.

The ***RAND FILL*** button, when clicked, erases the current cell population and fills it with randomly placed live cells.

The ***DENSITY*** knob to the left of these switches controls how many live cells are placed in the world. Its value is expressed as a percentage from 1% to 50%. If you set it to 25%, for example, and click the ***RAND FILL*** button, one quarter of all the squares in the world, chosen at random, will contain a live cell.

---

[2] If that doesn't help, have you played the arcade game Asteroids? Remember how the asteroids, and your ship, would appear on the left side of the screen after they slid off the right side, and vice versa? Same idea.

Stutter

The Stutter processor takes small slices of audio and plays them a number of times, producing buzzes, beeps, and glitchy buffer effects. This processor has three knobs, labeled ***REPEATS***, ***LENGTH***, and ***PITCH***.

The ***REPEATS*** knob sets the number of times that the slice of audio is repeated. The available values are powers of two, that is, two, four, eight, 16, etc. on up to 128. Since the duration of this effect is set by the ***SEQUENCE LIFESPAN***, the length of the repeated slice of audio is equal to the ***SEQUENCE LIFESPAN*** divided by the ***STUTTER REPEATS***. So for example if the ***SEQUENCE LIFESPAN*** is 1/8 and the ***REPEATS*** is four, the Stutter audio slice will be equal to the duration of a 1/32 note. That audio slice will be repeated four times when the Stutter processor is triggered.

Since the ***SEQUENCE LIFESPAN*** ranges from a quarter note down to a 1/32 note, and since the Stutter ***REPEATS*** parameter goes from 2 to 128, it follows that the Stutter processor uses audio slices ranging in length from that of an eighth note down to the equivalent of a 1/4096 note, if there were such a thing. This is why Stutter can sound like anything from a short loop of audio to a high-pitched tone. At a tempo of 120bpm, a 1/4096 note is only a little more than 1/10,000 of a second long. Our ears don't perceive something that happens 10,000 times per second as discrete events; instead we hear a high-pitched tone with a frequency of 10,000Hz. It is also entirely possible, depending on the signal you're running through Automaton, that there is nothing but silence within the first 1/4096 note of the audio present when the Stutter effect is triggered. This means that at some combinations of extreme settings, there's a good chance that you won't hear anything at all when Stutter is triggered.

The ***STUTTER LENGTH*** knob acts like a gate that silences part of each repetition of the audio slice. This knob can be set from 1% to 100%. At 100%, the gate is on all the time and has no effect. As you rotate the knob counter-clockwise, the gate starts to silence part of the audio slice. If the ***LENGTH*** knob is set to 50%, the second half of the slice is silenced each time it is played. Since the gate duration can be as short as 1%, you may not hear anything at all at low settings of the ***LENGTH*** knob, depending on your input material and the

length of audio slices that Stutter creates (which is determined by the ***STUTTER REPEATS*** knob and the ***SEQUENCE LIFESPAN*** knob, as described previously).

The ***STUTTER PITCH*** knob changes the pitch of the audio slice as it repeats. Each time the slice is played the rate at which it plays is increased or decreased slightly, causing its pitch to rise or fall with each repetition. If the knob is set to its center position the pitch of the audio slice is not changed. If you move the knob clockwise, the pitch rises; if you move the knob anti-clockwise, the pitch falls.

Note that the length of the audio slice can affect its perceived pitch. As described above, if the audio slice is short enough, the pitch you perceive will have nothing to do with the signal that Automaton is processing, and the Stutter Pitch knob will have no audible effect.

Modulate

The Modulate processor is a ring modulator. Ring modulators are familiar to many users of synthesizers and guitar effects; they produce "clangorous" or inharmonic tones by multiplying the input signal with the output of a built-in oscillator. Ring modulation is actually a form of amplitude modulation; "ring" refers to the arrangement of diodes in analog circuits originally used to create the effect. Automaton's Modulate processor includes a second oscillator which modulates the frequency of its main oscillator, creating a wider range of effects than possible with a simple ring modulator.

The Modulate processor has four knobs. The first knob, ***AMOUNT***, controls the amount that the incoming signal is modulated (affected) by the main oscillator. Turning this knob up increases the overall intensity of the effect.

The second knob, ***FREQ*** (short for frequency), controls the frequency of the main oscillator. Rotating this knob changes the timbre or tone of the effect.

The third and forth knobs control the second oscillator which modulates the frequency of the main oscillator. The **RATE** knob controls the rate of the second oscillator. You'll notice that this knob has a wide range, from one cycle every ten seconds (0.1Hz) up to 200Hz. At low frequencies you'll hear the modulation pitch rising and falling. At frequencies above 20Hz or so the frequency modulation will blur into a steady tone, and the overall effect of the Modulate processor will become more intense. The ***DEPTH*** knob controls how much the second oscillator modulates the main oscillator. You won't hear the effect of the second oscillator at all unless this knob is set to something higher than zero.

Bitcrush

The Bitcrush processor applies several different forms of digital signal destruction. This processor creates the popular "lo-fi" sample-rate reduction and bit-crusher digital effects we've all come to know and love despite having 24-bit audio converters and 64-bit signal-processing math readily available to us.

This processor has three knobs. The first knob, **RATE**, controls a sample-rate reducer. As you turn this knob up, the signal is resampled at a lower rate than your host's current sampling rate. If this knob is turned all the way anti-clockwise, the resampling process has no effect on the signal. As you turn the knob up, the signal is sampled at a lower rate. If you turn this knob all the way clockwise, the signal is sampled at 1/100th of your host's sampling frequency (e.g., 441Hz if you use the usual sampling frequency of 44100Hz).

The second knob, **BITS**, controls a bit-depth reduction or "bit-crushing" process. If you leave the knob turned fully anti-clockwise, the signal is passed with full resolution. (Automaton uses 32 bits to represent signals internally, but your host software may use 16, 24, or 32 bits.) As you rotate the knob, the number of bits used to represent the signal first drops to 16, and then decreases all the way to one as you turn the knob fully clockwise.

The third knob, **ERROR**, is an Audio Damage original. It introduces errors in the bits used to represent the signal. The knob controls how long the errors persist, and hence how much they damage the audio. If the knob is rotated fully anti-clockwise, no errors are added to the signal. As you rotate the knob, progressively

more persistent errors are added to the signal, creating digital-sounding noise. The **BITS** and **ERROR** knobs are somewhat complementary in that the effect of the **ERROR** function is less noticeable at high settings of the **BITS** knob.

Replicate

The Replicate processor is a scaled-down version of our wildly popular Replicant plug-in. In a manner similar to the Stutter processor, the Replicate processor records and repeats chunks of audio. Replicate operates on longer chunks than Stutter and hence is useful for creating repetitions of notes or short phrases and delay-like effects.

The **LENGTH** knob sets the length of the chunks of audio that the Replicate processor repeats. The setting of this knob is multiplied by the setting of the **SEQUENCER LIFESPAN** knob to determine the length of the chunk. The **LENGTH** knob has a range of 1/4 up to 16 (in powers of two), and the **SEQUENCER LIFESPAN** knob has a range of a 1/32 note up to a quarter note, so Replicate's chunks can range in duration from that of a 1/128 note up to a full measure. For example, if the **SEQUENCER LIFESPAN** knob is set to 1/16 and the **REPLICATE LENGTH** knob is set to 2, **REPLICATE** will repeat segments of audio equal in duration to an eighth note.

The **REPEATS** knob sets the number of times that the Replicate processor repeats a slice of audio. At its lowest setting of two, you'll hear each slice of audio played twice. The maximum setting is 16 repeats. This means that the Replicate effect can persist for quite awhile--up to 16 measures, if the **SEQUENCER LIFESPAN**, **REPLICATE LENGTH**, and **REPLICATE REPEATS** knobs are all set fully clockwise.

The Replicate processor includes a low-pass filter which is controlled by the **FILTER** knob. (The knob's name probably doesn't come as a surprise.) The frequency of the filter changes with each repetition of audio created by the Replicate processor. The first time the audio passes through the processor, the filter is set fully open so that it has no audible effect. With each successive repetition the filter closes, removing high frequencies and

causing the repeated audio to sound darker or duller. The amount that the filter closes is set with the **FILTER** knob. The setting of this knob determines the frequency of the filter used during the final repetition.

The **FILTER** knob has a range of 40Hz to 20kHz. At high settings the filter won't close much and you'll hear little or no effect. At low settings the filter's frequency may go below the lowest tones in Automaton's input signal and you may hear nothing at all after the first few repetitions.

The **DECAY** knob determines how much, if any, the repeated audio diminishes in loudness each time it is played. If the **DECAY** knob is rotated fully counter-clockwise, the looped audio is played back with the same loudness with each repetition. If you rotate the knob clockwise from this position, the looped audio becomes quieter each time it repeats. At the fully clockwise position, the audio usually fades out altogether before you hear the last repetition. You can use this control to make the Replicate processor sound more like a traditional delay effect. This control has a range of 0dB (no decay) to -6dB per repetition.

The Replicate processor's last control is the **TRIGGER** switch. If you've been paying careful attention, you might have realized that Replicate is the only processor in Automaton whose effect on the incoming audio can persist for longer than a single generation after receiving only one trigger. (This follows when you consider that the length of the audio segment that the Replicate processor records and repeats can be greater than the **SEQUENCE LIFESPAN** duration, since the **REPLICATE LENGTH** knob has a number of settings greater than one.) So, this question naturally arises: What happens if Replicate is triggered, starts merrily repeating some audio, and then receives another trigger before it finishes repeating? The answer is that it depends upon the setting of the **TRIGGER** switch. If the **TRIGGER** switch is set to **IGNORE**, the second trigger (and any subsequent triggers) is ignored until the Replicate processor finishes repeating the audio that it's currently playing. If the **TRIGGER** switch is set to **RETRIG**, Replicate will start recording and repeating a new segment of audio every time it receives a trigger.

Master

The Master pane has controls which affect Automaton's overall output level and how it mixes the processed audio it creates with the incoming audio.

The ***OUTPUT*** knob controls the loudness of Automaton's output signal. Use this knob to balance Automaton's overall output against other instruments in your mix. (Of course you can do this other ways in your DAW, but we thought it would be handy if Automaton had its own loudness control.) The ***OUTPUT*** knob has a range of -40dB (very quiet) to +6dB for some boost.

The ***FX LEVEL*** knob controls the loudness of only the processed signal. Use this knob to balance the audio that Automaton has mangled with the loudness of the original, unprocessed signal. Like the ***OUTPUT*** knob, the ***FX LEVEL*** knob has a range of -40 to +6dB.

The ***FX BYPASS*** switch bypasses all of Automaton's signal processors. Leave it set to ***OFF*** for normal operation, switch it to ***ON*** if you need to defeat Automaton's audible effect altogether. (Yes, like Automaton's other controls, you can automate this switch and/or control it with a MIDI controller. Have an appropriate amount of fun.)

The ***FX MODE*** buttons choose different signal-mixing configurations. The different configurations control how Automaton mixes the processed signal with the incoming signal to form its output signal, and are appropriate for the different ways in which a plug-in can be inserted in the signal paths in your host software.

The ***MIX*** mode mixes the input signal with the processed signal. This mode is appropriate when you’re using Automaton as an insert effect on one channel of your host software’s mixer, and you want to hear both the original signal and the processed signal mixed together.

The **DUCK** mode silences the input signal when it is playing back looped audio, and passes the input signal without alteration otherwise. This mode is useful when you're using Automaton as an insert effect, and you want to hear the original signal and the processed signal but not both at once.

The **SEND** mode does not pass any of the unprocessed input signal to Automaton's output. You hear only the looped audio. This mode is useful when you're using Automaton as a send effect in your host's mixer, sending one or more channels to it and mixing its output with the other signals.

These mixing modes have less obvious but interesting uses. For example, you can create a rhythmic gating effect by using the Send mode while Automaton is used as an insert effect.

MIDI Controllers

Automaton responds to MIDI continuous controller messages. You can use hardware MIDI controllers, such as MIDI slider boxes or the knobs found on some MIDI keyboards, to adjust Automaton's parameters. You can assign MIDI controllers to the knobs, sliders, and switches in Automaton's Sequence, effects, and Master panel. MIDI control is not available for the contents of the world grid.

Automaton has a simple "MIDI Learn" mode for assigning its controls to MIDI controllers. To assign a control to a MIDI controller:

- First, hold down the SHIFT and CTRL keys on your PC's keyboard, or SHIFT and CMD keys if you're using a Mac, and click once on the control. A orange box will be drawn around the control to indicate that it is ready to learn which MIDI controller it will be assigned to.

- Next, move the MIDI controller to send a continuous controller message—turn the knob, press the button, move the slider, whatever is appropriate.

- The orange box will disappear. Now the control will move when you manipulate the MIDI controller.

Automaton waits until it has received two consecutive continuous controller messages with the same controller number before it makes an assignment. This filters out extraneous data sent by some MIDI controllers. If you are assigning a button or switch on a MIDI controller, you may have to press or move the switch twice before Automaton recognizes the controller and assigns it to the desired control.

To assign a different MIDI controller to a control, repeat the same procedure using a different controller.

To cancel MIDI Learn mode without assigning a controller, hold down the SHIFT and CTRL keys (SHIFT and CMD keys on a Mac) and click in any empty area in Automaton's window (i.e., don't click on another control). The white box will disappear.

To remove a MIDI controller assignment from a control, SHIFT and CTRL keys, (SHIFT and CMD keys on a Mac) click on the control once so that the white box appears, then click again on the same control.

Automaton's MIDI controller assignments apply to all presets and instances of Automaton, in all host applications that you use. The MIDI assignments are stored in a special file on your hard drive. The contents of this file are read when Automaton is loaded by your host. If you have two or more instances of Automaton

in use at once, any MIDI assignments you make will not be propagated to the other instances until the next time that your host loads the plug-ins.

Automation

All of Automaton's parameters can be automated using your host's automation features; the contents of the world grid cannot be affected by automation. Consult your host's documentation for information on how to use these features.

And Finally...

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