

Elder Thing

a VST software Instrument by xoxos

Quid est?


The memorable description of Lovecraft's monster intrigued me as a target for waveguide synthesis. This product does not license any intellectual property and I am not particularly given to the genre, so connection between inspiration and product ends there... for purposes of public admission.

The description of the Elder Things (occasionally Old Ones) is more complex than this model, having in addition to a central mouth and five radiating breathing tubes (termed 'radii' herein), five additional mouths, and certainly no indication of internal dimensions. Interested in the musical properties of the multiport junction, I envisioned the radii adjoining at an intersection of the vocal tract.

At this point the model became strictly relevant to the source of inspiration. The radii were found to create resonances in a similar way to the nasal cavity's effect on the human vocal tract, being the primary channel of oscillation. In an effort to simplify the model for the sake of control, the radii was limited to two dimensions - length, and aperture at the junction.

In essence, the sum of the parts is a segmented Kelly-Lochbaum tract model coupled with a bank of resonators capable of adding a swarm-like effect to the central voice. This instrument is not well-suited for general Kelly-Lochbaum tasks. The complexity of the implementation is suitable for elementary articulations in the emulation of the hypothetical form.

Meeting the request for external audio to be routed through the tract as a resonant system, perhaps Elder Thing VST is in essence, a weird, creepy noise maker.



Technical Conception

Waveguide synthesis models the transmission of signals in physical space by using delays. A constant of 1130 f/s is used to dimension the atmospheric transmission of sound per the parameter settings to produce the desired emulation.

The conjectural vocal tract is divided as seen in the illustration, into two four-segment sections, one before the multiport junction, referenced as the trachea, and one following, referenced as the mouth. The term syrxinx was appropriated for the laryngeal folds to camp it up a bit.

It is advisable to be familiar with use of the Kelly-Lochbaum tract model before patching (excepting the sole intent to produce horrible noises) via the many free apps illustrating the circuit, or through use of my precedent Fauna VST. The length of the trachea and the length of the mouth can be separately indicated.

The Kelly-Lochbaum tract uses a set of ratio coefficients representing the cross-section area of each segment. Reflection between adjacent segments is computed, having the effect of adding formants to the signal generated by the waveguide. Ratio coefficients are also used to discern the signal reflection between the vocal tract and the radii.

The syrxinx is modeled with a second-order mass-spring, definable as all mass-springs in the instrument by resonant frequency (simply f on the plugin interface) and gain, or retention of energy. The syrxinx is functionally a reed valve, to open and close in response to the balance of pressure from the imagined lungs on one side, and returning pressure from the vocal tract on the other.

The overall pressure in the trachea drives a mass-spring which models tissue elasticity and modulates the length of the trachea. Similar mass-springs also modulate the length of the radii. The flex parameters are used to indicate responsivity of the trachea and radii. You may also wish to note that the coefficients for the trachea mass-spring are also used for the second order of the syrxinx mass-spring, supposing that it is mounted on tissue of similar consistency.

Signal from the tract is tapped at three points: the trachea, processed by a resonant lowpass filter to model sound passing through the body of the creature, the open mouth, and a stereo image produced by the five radii.

Instrumentalisation


I have established my production of VST with a discrete ethic of parameterisation, and Elder Thing follows course. If it's a coefficient, there's a parameter for it on the GUI. While this may seem bewildering at first, I would rather have too many parameters than too few. With moderate application I am confident that the perspective of the user will change to consider this instrument elementary and simplified. There were several parameters that I excused with the objective that at some point, any instrument that is used needs to be completed and released.

The advanced user will consider that it would be wonderful to have control over the reflection, formantisation, et c. of each radii. The initial user will appreciate that simplification of the parameter set to a naive "scare" and "scare more" setting would offer little in the way of musical flexibility. I cannot effectively propose a gestural solution et al. for the control of this instrument as I and other users do not have a consistent accoutrement beyond the DAW.

To improve the circumstance, a 13 slot send panel is provided to augment the control environment. Internal and external modulators can be assigned to multiple destinations, including the amount parameter of subordinate send slots.

The primary dimension of expression is a bus which can distribute modulation to every tract parameter in the upper panel by means of a strip of rotary trimmers under each slider. Note that similar devices associated with the amplification and modulation envelopes instead apply MIDI note velocity.

There are four sets of coefficients for the trachea and mouth reflection ratios which may be cross-faded using respective x-y "joystick" pads. Because switching patches has no effect on the model beyond altering parameter values, configuring several presets can be used to allow x-y pads to crossfade between more values.



Modulation Guidelines

Destinations are, with minor exception, scaled from 0 to 1.

Envelopes and external modulations apart from pitch bend are ranged from 0 to 1.

The four assignable LFOs, and pitch bend are scaled from -1 to 1 (with the exception of the random contour), indicating that a median setting will sweep the range of most destinations.

There are two specialised radii LFOs which generate five outputs and may only be assigned to radii length or ratio. These LFOs have a total range of 1 (eg. from -.5 to .5).

The x-y pads are scaled from -.5 to .5.

Most parameter targets have an upper and lower limit where sensible. Some parameters may be modulated beyond the range of the GUI setting, for instance the pressure parameter can receive negative modulation beyond the slider's lowest setting of 0.

Send Panel

The send panel allows modulators to be assigned to several destinations in varying amounts. Each send slot has a source, destination, amount trimmer and invert button. Amount trimmers have a linear scale. Since slots can be assigned to the amount of slots above them, two slots may be used to produce a curved response from a linear source, by using the same source for both slots and the second to modulate the amount of the first.

Note that sends may not use modulators as destinations, apart from the five output radii LFOs. Envelopes and LFOs may be assigned to LFOs outside of the send panel.

In order to accomodate hosts that do not have sample accurate automation, there are four smoothed parameters. These are not visible on the GUI but appear in the host's automation targets as asmooth 1-4 (the 'a' is to prioritise them alphabetically). External modulation applied to these targets is averaged across 256 samples. Smooth 1-4 may then be assigned to any destination. Use of cc#s 1-4 may provide improved performance in some cases as well.

X-Y Pads

The mouth x-y pad can be locked to the trachea x-y pad for convenience.

LFOs

LFOs may be synced to NoteOn. They may not be synced to host tempo because Euclid would approve.




(radii LFO only)

There are seven available contours (eight on the radii LFOs) and a CPU-saving off position. The three linear contours can be crossfaded between saw, triangle and ramp waveforms using the width parameter, coded in red on the GUI. All contours except the first linear contour may also be altered by the mod parameter. The modification is illustrated in gray.

The random contour uses cubic interpolation between periodically generated values to produce a varying, sinelike contour. Because of the interpolation method, it's possible that LFO peaks may slightly exceed (-1, 1). Moreover, the mod contour for the random LFO is similar to the first sine waveform - you may recognise this contour as $(n * (abs)n)$. Because of the squaring, peaks may occasionally exceed (-1, 1) by a higher amount when mod is used. This setting rarifies variation from the center making peaks more pronounced.

LFOs (continued)



Rarefaction of variance in randomised (and sine) LFO waveform with mod setting due to $n * (\text{abs})n$

Random contour LFOs do not respond to the phase parameter except as a seed value when the LFO is reset.

The radii LFOs produce five signals which may be spread from having the same phase to being evenly distributed across the wavecycle. This produces two options for the random contour option - to have the five outputs follow the same sequence of generated values, with spread serving to increase follow time between the outputs, or the rand 5 option where each output has its own individual and unrelated output.

Envelopes

Elder Thing VST makes use of a 16 stage graphic envelope resource created by Chris Kerry. Clicking anywhere on them outside of the boxes produces a menu where the number of stages, contour of each stage, sustain and repeat options are available. You may also change the size of the envelope editing boxes for ease of use.

Outside of the graphic interface are buttons to select looping in release (marked *lir*), inversion, rate, amplitude and velocity sensitivity for rate and amplitude. Note that the rate of the envelope stages may be individually adjusted to great effect.

Calibration

Elder Thing VST parameter scales are not indicated on the GUI, again in order to balk and confound Euclid. Should you wish to tailor the model to a more substantial universe, the following informations may be of use:

noise lp:	86.13 Hz to 22050 Hz (12 dB / octave)	\wedge^2 scale
noise hp:	13.75 Hz to 14080 Hz (12 dB / octave)	\wedge^2 scale
trachea length:	0.5 ft to 5 ft (longer than man)	linear scale
mouth length:	0.25 ft to 2.5 ft	linear scale
all mass-springs:	13.75 Hz to 14080 Hz	\wedge^2 scale
trachea flexibility:	0 to 50% of total length	linear scale
radii flexibility:	0 to 10% of total length	linear scale
damping:	13.75 Hz to 14080 Hz (6 db / octave)	\wedge^2 scale
transduction:	0 to 1.0	linear scale
trachea mix f:	21.53 Hz to 22050 Hz (12 dB / octave)	\wedge^2 scale
radii length:	0.5 ft to 8 ft	linear scale
all lfo rates:	297.89 s to 110 Hz	\wedge^2 scale

My Preset Sounds Different?

The audio samples in the waveguide are not zeroed out between MIDI notes or patch selection, and, like any delay system, the timbre will depend on the summing of current and new samples. The timbre of a less volatile patch may differ if switching from a volatile patch (an example of this is the drone patch, which may produce a keening timbre instead of a deep, FM sound if the tract contains a previous signal).

This can be avoided by using an envelope to raise input parameters (eg. pressure) from zero on each note.

Please test demo version to ensure suitability before purchase. There is no expression of guarantee. The words stygian, cyclopean and non-euclidean are not used in this document. Elder Thing VST was created with the SynthEdit SDK - www.SynthEdit.com

Thanks to Chaosium Inc. for indicating that they hold no property claim on the phrase Elder Thing

©2011 2012 Rurik Leffanta - All Rights Reserved - VST is a trademark of Steinberg Soft und Hardware GmbH.