

Water VST

stochastic liquid sounds synthesizer by xoxos

Installation

Copy the water.dll file to your VSTPlugIns folder.

Algorithm

An air bubble introduced to water acts like a mass-spring, oscillating as the pressure difference between the liquid and bubble reaches equilibrium. For spherical bubbles, this oscillation is at a single frequency determined by the radius, as is the decay rate (or 'damping' ;) and amplitude of the oscillation. Nonspherical bubbles exhibit modes of oscillation in the same way plates, membranes and instrument bodies exhibit vibrational modes. For the sake of efficiency, this algorithm assumes spherical bubbles, and that a droplets interaction with the surface of water can be handled as a bubble.

The impulse response of the bubble is accordingly a decaying sine, which is very convenient for synthesis. As is easily surmised from experience, the frequency of oscillation also rises as the bubble approaches the surface and pressure is reduced. Kees van den Doel notes in Physically Based Models for Liquid Sounds that a bubble at the surface has a frequency $\sqrt{2}$ higher than a submerged bubble (which he considers to be pragmatically depths greater than 10 times the radius). Greater pitch variation can result from other factors.

In most parts, parameterisation for this model is physically calibrated, so that no guesswork is required for emulative patching. Pitch rise as an exception to this would require more deterministic variables than the more musically useful user specification.

Water VST is then a simple application of parameterised stochastic method to the generation of multiple exponentially decaying sinusoids with optioned pitch rise. Gaussian, or bell-shaped distribution curves are applied to the randomised coefficients to produce natural sounding variation.

As everyone is aware, even the regular drip of a faucet results in continuous variation. Water VST includes a secondary event circuit so that even rigidly clocked single drips occasionally produce a second droplet sound which can be virtually simultaneous or lag behind the initial event by up to 80 milliseconds.



Parameterisation (Top Row)

This section regarding the top row can be ignored as it should be fairly intuitive and straightforward in use. The next section regarding the bottom row should be read in order to best use the implementation.

Radius & Radius Bandwidth

The frequency in Hertz of a spherical bubble is equal to $3 / \text{radius (meters)}$, so that the acoustic range of 20 Hz to 20 kHz is equivalent to radii of 15 cm to 0.15 mm. The radius slider has a maximum radius of 2cm, or 150 Hz. Pitch modulation can be applied to produce significantly lower frequencies.

The radius bandwidth parameter increases the range of variation between the radius setting and each event. The maximum setting corresponds to a range of 0.1 to 10. All bandwidth parameters use gaussian distribution, which rarifies probabilities at the extremes of the range as corresponds to most cases of natural variation.

Depth & Depth Bandwidth

Depth indicates the depth of pitch modulation. The center position corresponds to the $\sqrt{2}$ coefficient mentioned earlier, and the maximum setting is twice this (over an octave). Depth bandwidth increases variation around this center coefficient.

Pitch

The pitch control indicates the rate that the pitch rises. At the lowest slider position, the frequency of the oscillation will reach the setting indicated by the radius slider at the end of its acoustic decay. Higher settings correspond to faster rates. The highest setting tends to drive the pitch of every event to nyquist by the end of its amplification.

Rate & Rate Bandwidth

The rate parameter sets the periodicity of bubble events. The highest setting corresponds to two seconds in between events, the lowest, fastest setting to two samples. The bandwidth varies the event clock from perfectly regular to zero to twice the period.

Bright

The amplitude of each bubble corresponds to its size and is, without consideration of the properties of the liquid, equal to the radius times the square root of the radius.

Normal amplitude corresponds with the center position of the slider. Higher settings are appropriate for close sounds, and lower settings are appropriate for distant sounds.

Wide

This parameter sets the stereo width of the signal. The lowest position is mono, and the central position corresponds to full stereo width. Higher settings attenuate the volume of hard panned sounds to increase the perceived stereo width.

Field

Field increases the amplitude variation in events. The lowest setting amplifies all events relative to their radius, increasing the setting will lower the volume of some events, creating the perception that they are occurring at a distance (eg. rainfall would use a high field setting, whereas a running tap would use a low field setting).



Parameterisation (Bottom Row)

Read this section! :)

Mod

Mod applies modulation to a group of parameters simultaneously. Knobs under the top row of sliders trim the amount of group modulation applied, so that eg. a trickling brook can be gradually transformed into a waterfall.

There are four sources of modulation that are applied to these trimmers: the mod slider, the modulation wheel, pitch bend, and the modulation envelope.

For most parameters, modulation is applied in the direction associated with increased water flow or energy. Some parameters include an invert button under the trimmer in case the opposite effect is required.

Modulation Envelope

This ADSR envelope features pickup behaviour, like the amp envelope, so that for example, cresting waves can be played using the keyboard.

There are two rows of knobs under the envelope. The first, as is common to all of my plugins, trims keyboard velocity. The two knobs under the rate parameters pertain to attack, and decay/release rate. Higher velocity results in faster attack and longer decay and release.

The second row of knobs indicates the curve of the attack and decay/release contours. The central position indicates a linear contour. The lowest position indicates a concave slope (50% time corresponds to 25% amplitude) and the highest position to a convex slope (50% time corresponds to 75% amplitude). A conventional "analog" envelope would have the attack curve in the highest position and the decay/release curve in the lowest position.

FYI the amplitude envelope uses a linear contour on the attack portion (green) and an exponential, convex natural sounding release (red).

Sync & Seed

Sync produces recallable performance on key triggers. Because there is no key tracking or pitch, each key has a different recallable sequence. Changing the random seed parameter will alter the sequence for each key.

The timbre may not be perfectly identical on each trigger as the sync feature only resets the value used to generate the sequence. If the algorithm is generating a "secondary droplet event" when the new midi note is received, the reset seed value will be applied to the creation of droplet event, which uses a different number of random functions.

CPU

This number indicates the total number of simultaneous 'droplets' or events. This parameter allows for some creative and important performance features, eg. if the clock rate is faster than the decay of a droplet event and cpu is set to 1, no new event will occur on clock pulses until the current droplet has ended, producing rests in the rhythm.

The cpu parameter acts to set the sound density (as does effectively the field parameter).

The actual amount of cpu used also heavily depends on the radius parameter. Larger radii have longer decays and continuously use cpu, whereas small radii produce very short events. Much higher densities can be used with small radii using less cpu than large radii (where you can max out a cpu core with only a few hundred).

Be aware that the highest number of simultaneous events is 8192, so that suddenly raising this parameter in use may cause your cpu to pause, depending on how awesome your DAW is.

Amp

Be aware that low frequencies are considerably louder than high frequencies, so this slider has a very wide range! Use low values for deep pitched timbres and save new patches accordingly.

Patching Tips


The relative amplitude of each drop can be adjusted as a function of frequency, creating the perception of proximity to the source. Otherwise, there are no improvements to the synthesis in regards to perception. Some consideration is appropriate here in use for all emulative applications.

A location processing effect should be added to the signal chain if one is available to you. Otherwise, a multiband equalizer will help add some perceptual quality. Reverberation using allpass filtering may also be useful in adding a perception of the submerged region the water is flowing in in addition to normal use of reverberation to create a sense of space.

RS-MET offers a free graphic multiband equaliser available here - <http://www.rs-met.com/freebies.html>

An emulative setting is also likely to require multiple instances of Water VST in order to more precisely patch each dynamic present, eg. water flowing out of a faucet into a sink can easily be divided into water exiting the faucet, circulating in the basin and in the drain (which may benefit from a comb and lowpass filter or other reverb process). A scene featuring rainfall will sound more realistic if multiple patches are used in ensemble to build a scene depicting a variety of surfaces.

For host synced droplet sounds, set the rate slider to the highest position (slowest rate, or 2 seconds) and sequence different midi notes for each drip, or the same midi note and modulate the seed parameter for variation.



envelope curves

