AriesVerb

version 0.4a
Christian Schüler

December 1, 2008

AriesVerb is a versatile effect processor—it can do reverb, delay, echo, flanger, chorus, phaser, comb filters and body resonances, as well as combinations of all of the above. This is possible because AriesVerb is built on a Feedback Delay Network, which is a well known algorithm unifying all of the aforementioned algorithms into one framework. AriesVerb exposes the FDN in a novel way which results in musically meaningful controls. This document explains the principles behind AriesVerb and the controls exposed in the current version.

Part I
Function

Basics

AriesVerb is described most accurately as a vectorized comb filter. Conceptually it looks like a comb filter (a delay with feedback), but all elements are multiplied by four and a 4×4 matrix sits in place of the simple feedback gain. Such a structure is known as a *Feedback Delay Network*, and the complete schematic is shown in figure 1.

The signal flow is as follows: Sound x_k enters via the input vector b, which distributes the signal among the delay lines $z_1 \ldots z_4$. The delay lines hold the

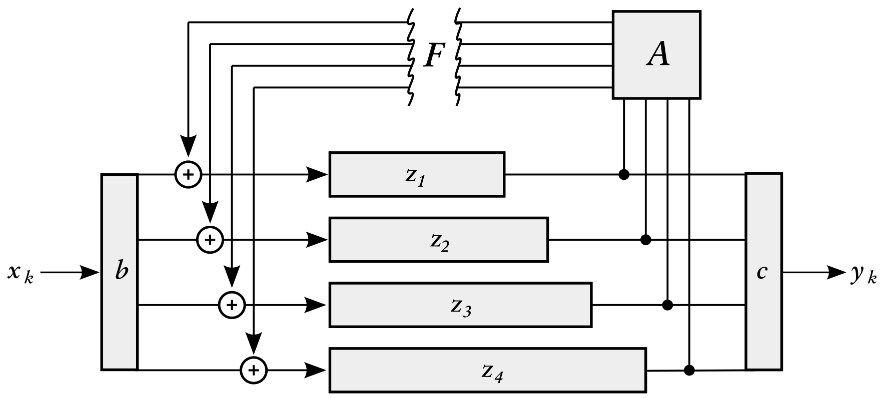

Figure 1: FDN structure of AriesVerb. A: Feedback matrix, b: input vector, c: output vector, x: input signal series, y_k: output signal series, z_i: delay line i of length t_i, F: filters and saturation.

signal for individual times $t_1 \dots t_4$. Each delay line can feed back into each other delay line with an amount specified in the feedback matrix A. Finally, the output vector c collects a weighted sum from the delay lines and forms the output signal y_k. In addition, there are filters and a non-linear saturation in the feedback loop (at position F).

Echo– vs mode density

AriesVerb has the potential to build a high echo density in a short time, which is evident from its FDN structure. A single impulse generates 4 echos, these generate 16 echos, these in turn 64, then 256, etc. However, a quality reverb needs both, a high echo density and a high mode density. The mode density is the density of the resonant peaks (modes) in the spectrum. A low mode density results in a metallic sound, while a low echo density results in a machine-gun effect. See figure 2b for an analysis of AriesVerb's default program Temple of the Ancestors.

How to control echo- and mode density is explained in more detail in the Parameters section. Unfortunately, echo- and mode density are complementary to each other, and increasing one will always decrease the other and vice versa.

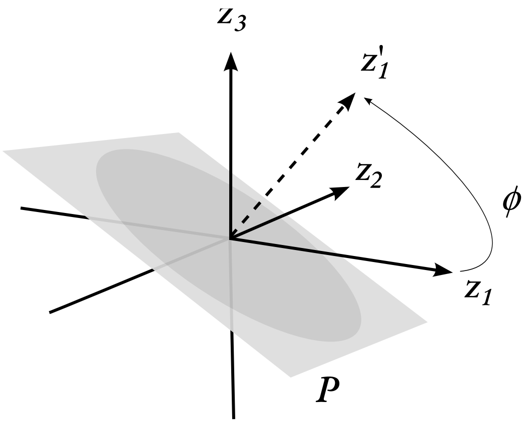

(a) Feedback rotation.

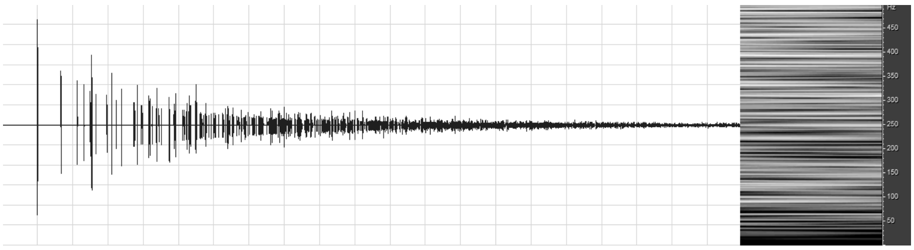

(b) Echo- and mode density.

Figure 2: (a): Feedback rotation. P: rotation plane, ϕ: rotation angle, z_i: signal at delay line i. z_1': image of . (b): Echo- and mode density. Left: Impulse response over 1 second time, right: a high resolution spectrum of the same signal.

Feedback rotation

The job of the feedback matrix is to intermix sound between delay lines while preserving the overall energy. This is accomplished by means of a *rotation matrix*, as illustrated in figure 2a.

To understand this concept, think of the matrix action as a rotation of a 4D-vector into another 4D-vector. The rotation takes place parallel to a *rotation plane* with an amount specified by a *rotation angle*. The example rotation in figure 2a shows how a signal coming from delay line 1 (represented by the z_1-axis) would be rotated into a position representing an equal mix between delay lines 2 and 3 (stippled line). A 4×4 matrix can actually rotate in two planes simultaneously, which is the reason AriesVerb has controls for two independent rotation angles. Unfortunately such a 4D-rotation is impossible to illustrate on paper. There is more on this topic in the Parameters section.

Input– and output vectors

At the input stage of the FDN the sound is divided into 4 portions and assigned to the delay lines by amounts specified in the *input vector*, and at the output stage the sound is summed from individual delay lines by amounts specified in the *output vector*.

In AriesVerb 0.4, the direction of input- and output vectors (and therewith the share given to each delay line) is specified with only a single angle. Figure 3 shows the currently implemented layout. For stereo, the vector for the right channel is rotated 90° clockwise from the corresponding vector for the left channel.

Part II
Parameters

This part documents all parameters exposed by the AriesVerb plug in. As of version 0.4, AriesVerb has 16 parameters plus 2 meta-parameters grouped into 5 categories. The parameter names have been truncated to comply with the 8-character limit mandated in the VST SDK.

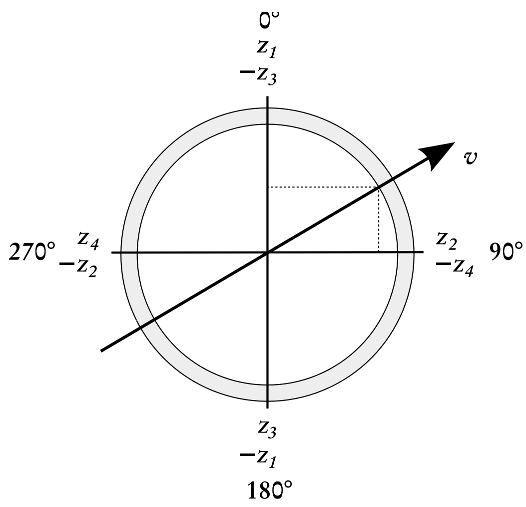

Figure 3: An input- or output vector (v) on the distribution plane with perpendiculars casted onto the axes representing delay lines $z_1 \dots z_4$.

Delay body parameters

Geometry: Selects one out of 17 pre-defined delay length ratios. They are documented in table 1.

Base dly: Sets the length t_1, the shortest of the delay lines, in seconds.

LF life: Half-life for the low end of the spectrum, in multiples of t_1.

HF life: Half-life for the high end of the spectrum, in multiples of t_1.

The geometry and base delay parameters together control the lengths $t_1 \dots t_4$ of the delay lines. The geometry parameter selects a fixed ratio for all delay lines while the base delay scales them (However, the maximum time for any delay line is 1 second).

The half-life parameters control the feedback strength. Here, *half-life* is the time by which a sound in the feedback loop decays to *half* its amplitude (-6 dB). Therefore, ten times of it correspond to RT60. The half-life is not specified directly, but in terms of a factor relative to t_1. This way, the half-life has a musical meaning of “resonance”, equivalent to the Q of a filter. Feedback is disabled altogether if any half-life is zero.

Name	t_2/t_1	t_3/t_1	t_4/t_1
SPHERE	1	1	1
BOX$_1$	1.04	1.11	1.16
BOX$_2$	1.05	1.14	1.3
DENSE$_1$	1.22	1.34	1.72
DENSE$_2$	1.78	2.44	2.62
DRUM	1.15	1.63	2.41
HALLWAY	1.41	2.22	3.13
ROOM$_1$	1.67	2.52	3.11
ROOM$_2$	1.44	2.59	3.35
SPARSE$_1$	2.24	3.78	4.67
SPARSE$_2$	2.61	3.78	5.1
MINOR	4/3	5/6	2
MAJOR	4/3	8/5	2
HARMONIC	4/3	2	4
OCTAVES	2	4	8
FIFTHS	3/2	2	3
DETUNED	+0.01%	+0.3%	+0.31%

Table 1: Geometry presets.

Feedback parameters

FB PLANE: Selects one out of 6 pre-defined orientations for the feedback rotation plane.

FB ROT 1: First feedback rotation angle, in degrees.

FB ROT 2: Second feedback rotation angle, in degrees.

FB INV: Post-rotation inverter, spectral coverage in percent.

The feedback plane controls the type of matrix: The "dense" presets generate very balanced matrices for smooth reverb and high echo density. The "sparse" presets generate unbalanced matrices with more audible graininess. The "block" presets generate 2×2 block matrices.

The rotation angles control the overall amount of coupling between delay lines. The effect of different rotation angles on the generated matrices is shown in table 2:

- With no rotation (both angles are zero) there is no matrix action and the system reduces to an array of parallel comb filters.

- The coupling is maximal if one rotation angle is 90° and the other is either 0° or 180°. This is the best setting for reverb.
- With block matrices, the system works as a pair of FDN pairs. The first rotation angle controls coupling between delay lines 1 and 2, and the second rotation angle controls coupling between delay lines 3 and 4.

Finally, the last parameter in the feedback category controls the spectral coverage of a frequency-dependent phase inverter. This inverter is neutral at 0% and inverts the full spectrum at 100%. At intermediate settings the inverter will only invert part of the spectrum. The inverter is still available if feedback is turned off, it is then placed between delay lines and output.

Non-linear effects

Mod dep: Depth of the delay length modulation, in percent.

Mod per: Period of the delay length modulation, in seconds.

Mod spr: Spread (phase difference) of the delay length modulation, as percentage of complete equidistribution.

Amp drv: Shape of a non-linear saturation function in the feedback loop, in arbitrary units.

The modulation parameters control how the delay lengths change over time. The length of each delay line will oscillate about the base delay length by a percentage specified as the modulation depth, and with frequency which is the inverse of the the modulation period. The continuous dilation and contraction of the delay lines causes a pitch shift similar to a Doppler effect, and is proportional to the absolute rate of change in length. Typical FM effects are introduced if the modulation frequency is in the audio range. If the modulation spread is zero, all delay lines contract and dilate in unison. At full spread, there is always one delay line moving in the opposite direction of another, so the average pitch shift is zero.

The Amp drv parameter controls the shape of a non-linear saturation function very similar to tape-, tube- or transistor saturation (without hysteresis). Example shapes are graphed in figure 4 for a control parameter in the range $1/2 \ldots 2$. The saturation is sandwitched between a pre- and a post-filter to reduce aliasing. An RMS correction at the output is optimized for a -12dB(FS)[1]

[1] Digital signals are measured in *decibels of the full scale*, dB(FS). A signal of 0 dB(FS) is at the threshold of clipping, and -12dB(FS) are 25% of this.

$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$ 0° | 0°	$\begin{pmatrix} 0.5 & -0.5 & 0.5 & 0.5 \\ 0.5 & 0.5 & 0.5 & -0.5 \\ -0.5 & 0.5 & 0.5 & 0.5 \\ 0.5 & 0.5 & -0.5 & 0.5 \end{pmatrix}$ 90° | 0°	$\begin{pmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix}$ 90° | 90°
$\begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$ 180° | 0°	$\begin{pmatrix} -0.5 & -0.5 & -0.5 & 0.5 \\ 0.5 & -0.5 & 0.5 & 0.5 \\ 0.5 & 0.5 & -0.5 & 0.5 \\ 0.5 & -0.5 & -0.5 & -0.5 \end{pmatrix}$ 180° | 90°	$\begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$ 180° | 180°
$\begin{pmatrix} 0.15 & -0.35 & 0.35 & 0.85 \\ 0.35 & 0.85 & -0.15 & 0.35 \\ 0.35 & 0.15 & 0.85 & -0.35 \\ -0.85 & 0.35 & 0.35 & 0.15 \end{pmatrix}$ 90° | 0°	$\begin{pmatrix} -0.71 & 0 & 0.71 & 0 \\ 0 & 0.71 & 0 & 0.71 \\ 0.71 & 0 & 0.71 & 0 \\ 0 & 0.71 & 0 & -0.71 \end{pmatrix}$ 180° | 0°	$\begin{pmatrix} 0.71 & -0.71 & 0 & 0 \\ 0.71 & 0.71 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$ 45° | 0°

Table 2: Example matrices for different rotation angles, generated with the DENSE$_1$ preset (top row and middle rows), the SPARSE$_1$ preset (bottom row) and the BLOCK$_1$ preset (bottom row, rightmost).

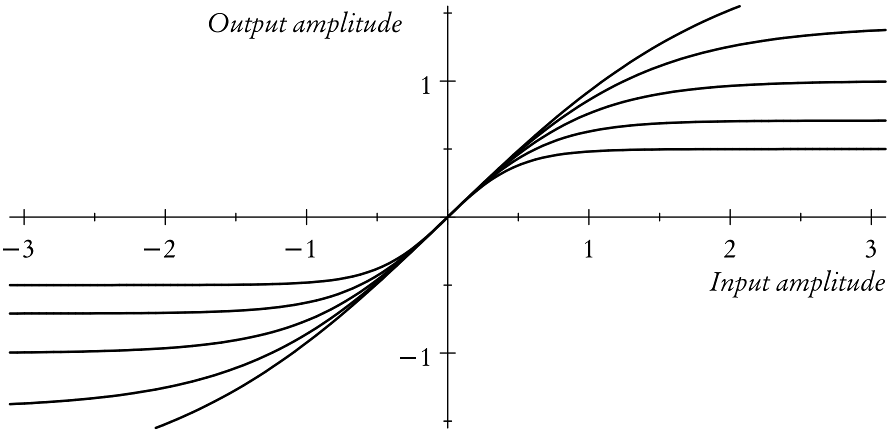

Figure 4: Saturation Function

signal, such that the perceived loudness of a -12dB(FS) signal is not altered regardless of the shape of the saturation function.

Input and output

In dir: Direction of the input vector, in degrees.

Out dir: Direction of the output vector, in degrees.

Dry mix: Amount of dry (unprocessed) signal, in procent.

Wet mix: Amount of wet (processed) signal, in percent.

The first two parameters control the direction of the input- and output vectors on the distribution plane, which have already been discussed (see figure 3 from the previous part). The last two parameters control the mix of dry and wet signal. If the input vector has the opposite direction of the output vector (180° difference), then the wet signal would be inverted.

Meta switches

Hold: Enables the lossless feedback loop, binary switch.

Dr. Who: Disables delay line reset on program change, binary switch.

These two parameters are "meta" parameters because they are not stored with any program; instead they change the behaviour about the programs[2]. The HOLD parameter enables the *lossless prototype*. In this mode, the input is muted and the remaining sound recycles endlessly through the feedback loop. If modulation is activated then this mode is not exactly lossless, due to slight interpolation error. **In case of strong, FM-like modulation the HOLD mode may cause a catastrophic build-up** (safely clipped, but beware of your ears).

The DR. WHO parameter disables the reset of the delay lines on a program change. If activated and AriesVerb resumes work with a new program, the delay lines are filled with whatever random sound was left from the last operation. This creates unpredictable clicks and pops that cycle further through the feedback loop. It may sound like some special effects from early science fiction movies.

Part III
Presets

AriesVerb 0.4 comes with 64 presets grouped into 7 categories. This part discusses the design and special aspects of each preset.

Environmental reverb

Environmental reverb is the domain for which AriesVerb has originally been developed. This is about creating the plausible illusion of a natural space with minimal CPU consumption. The 5 presets in this category are ordered by decreasing size of the resonant space, from very big to very small. They mostly use a dense matrix near 90° and a little amount of modulation to create a smooth irregular reverb. The quality of the environments can be increased when used together with early reflections (see Echos and Reflections).

TEMPLE OF THE ANCTESTORS: Simulating a large bright space like a stone hall with only little damping. The rotation angles are a little off (87°|5°) to create unsymmetry in the feedback matrix. The inverter is also a little off (99.7%) to create dispersion.

[2]This is not compatible with all known hosts. Some hosts try and implement program change "on foot" by querying and re-setting all parameters (instead of just sending a program change signal). Then of course, the HOLD and DR WHO parameters get overwritten, which is not the intended behaviour.

Open Stage in the Woods: Very similar to the above preset, but creating a slapback effect using geometry DENSE_1 with a large base delay. This geometry has all delay lengths between 1 and 2 times, so there comes a burst of echos after an initial pause the length of the base delay.

U Boat Club Party: Simulating a damp cellar environment. The high frequency half life is lowered significantly to create damping and the base delay is shorter.

Cardio Fitness: A smaller bright space with a significant amount of modulation, to simulate the thermal convections in the air between warm bodies and cold walls. No, really.

Very Small Dorm Room: A simulation of the author's living conditions from past days. The base delay is less than 10 milliseconds (3 meters) and there is no modulation.

Body resonances

Linear resonances from piano- or violin bodies can be seen as continuation of environmental reverb to a smaller scale. The body resonances in this category are chosen because they are interesting in some other aspect besides being "small rooms".

Biohazard Container: A simulation of the inside of a tin canister. This environment has a strong resonance (the half-life is high) with a strong damping.

It is Spring: The spring body simulates the humming of excited springs. The BOX_2 geometry has all delay lengths close together creating a regular, but not quite regular echo pattern. The feedback rotation is used to introduce low-frequency beating. Real springs would also have extremely high dispersion, which AriesVerb cannot simulate.

Oversized Shoe Box: A small body in the range of centimeters. At this level, the presence of the inverter is becomes relevant. The inverter simulates the reflection of the wave at an open terminal, creating the "hollow" impression.

M. . . ooh! : An even smaller body simulating the inside of a cowbell. It is not much different from the tin canister environment, except for the much shorter base delay. We are now in the microseconds (μs) domain showcasing the well engineered fractional delay lines of AriesVerb!

Glass Dish (fragile): A simulation of the resonance from very stiff bodies, such as glass. Geometry and feedback are chosen to create distinct and very sharp resonant peaks closely spaced around 8 kHz. A little amount of non-linearity simulates the coupling between longitudinal and transversal waves (also known as "clank").

Modulation effects

This category covers a wide range of effects for which delay length modulation is the central element. Phaser, flanger, chorus, et al. are here.

Trance I+II: The trance presets are simply large environments with strong modulation. The appearance is over-reverberated and thick. A physical analogy could be the reverb inside a large cavern with turbulent, moving air.

Airplaaaine: The airplane presets implement a basic "phaser" effect, causing a destructive interference to move through the spectrum as differently delayed signals meet at the output stage.

Stereo Airplaaaine: The same as above, but with a different input- and output vector.

Laurel & Hardy: A simulation of the trademark hollow-blowpipe sound used in slapstick comedy films. It is similar to the previous ones but has a strong feedback with inverter. This is a "flanger" type effect.

Howling Spirits: The same as above without inverter. Caveat: A high resonance without inverter produces a pole at DC (0 Hz) like in this preset. The result is bass overflow.

Instant Honk: This is a typical "chorus" effect which creates many individually detuned voices via pitch shift. It sounds like the detuned strings in a honky tonk piano.

PWM-alyzer: This effect mimics the trademark pulse-width-modulation sound. It is a combination of destructive interference (inverter) and a chorus effect.

Rotating Speakers: This effect is also a chorus but with a more synchronous modulation and a higher modulation frequency. The result is an audible vibrato as if the sound was played through rotating speakers, reminescent of 60's electrical organs.

A very Battered Tape: A single pitch shift with a little bit of feedback and a portion saturation, giving the impression of an aging tape recorder.

Auto Scratcher and Chaos Scratcher: These effects take pitch shift to the extreme, like a DJ was scratching the sound from vinyl (incl. reversals). The chaos scratcher also deliberately overdraws the delay lengths where they stay clamped at 1 second.

Distortion- and FM effects

This category covers effects which either use saturation to the point of distortion or audio-rate modulation. The result of both is the synthesis of new spectral content (harmonics and side bands).

Bass Fatten I. . . III: This is a chorus with distortion. The distortion creates additional harmonics on single notes, while the chorus increases the perceived bandwidth of individual harmonics. If used on a monophonic bass, it adds “fat”.

Tube Warmth/Presence/Scream: This is distortion without modulation, but with a strong coloring. A very short delay is employed here to simply work as an equalizer (“static phaser”), and different geometries create different spectral response curves.

Maximum overdrive: Very strong distortion with a bit of resonance, everything else is neutral.

Psychotic Padded Room: This is a medium sized environment with a very fast modulation frequency (500 Hz). This is equivalent to FM which creates side bands around each frequency at intervals of the modulation frequency. These side bands then merge with the reverbreation and become themselves sources for side bands, and so forth. Therefore, this room can respond at different frequencies than the ones that are put in, which *does* qualify as insane.

Washing Program 60° and 95°: Taking FM with feedback a step further creates noise. All the side bands that are recursively created from the input signal add up as one incoherent warbling.

Major Glitch: This is a pure FM effect and almost nothing else, and sounds like an artificially low-quality digital signal. The side bands from the modulation mirror the input signal at fixed intervals over the spectrum, very similar to what would happen in case of severe aliasing.

Insert effects

The presets in this category are examples of insert effects, which are designed to be applied to specific musical content. They are not general-purpose effects.

Drum Room XS/S/M/L: These effects are designed to add presence to a dry drum recording, and are prime examples of reverberation with high echo density, but low mode density. Drums are short, percussive signals with a broad spectrum. The delay lengths must be short enough to avoid audible flutter. The reverberation time can only be increased with more resonance. This however would make the low mode density apparent with metallic ringing. The solution on the medium and long presets is therefore to mask ringing with a significant amount of modulation and distortion, which works ok for drums.

Keys Insert XS/S/M/L: These effects are designed to add presence to dry piano or acoustic guitar recordings. They have higher mode density and lower echo density than the previous presets. While piano sounds are still somewhat percussive, the delay lengths must be long enough to avoid audible coloration. There is an upper bound on the percussiveness of sounds that can be used with these effects. Modulation is only used below the threshold of audible pitch shift.

Voice Insert S/L: These effects are designed to add presence to to purely tonal sounds like voices. They have very long delay lengths. The output vector is directed such that the left channel picks the first reflection.

Slapback Insert S/L: The difference of these effects with the previous ones, besides even longer delay lengths, is the complete lack of feedback. Therefore these effects are of the "one-shot" slapback type, and can be used to add some pop-style quality to voice recordings.

Echos and reflections

The presets in this category are examples of echo effects with the intention of audible, discrete echos. The echos must not merge into a smooth reverb tail, therefore, all delay lengths must have a simple common divisor.

Straight Echo: The simplest case is a regular echo. All delay lengths are equal (geometry SPHERE). Input- and output vector are equal, so left and right channel stay separated.

Stereo Echo: This preset produces different echo patterns on the left and right channel. All delay legths are multiples of $^1/_2 t_1$ (geometry FIFTHS). Feedback rotation creates a quasi-random echo pattern. The output vector is set to 45° to make the left and right channels a combination of all delay lines (with individual signs).

Rotating Echo: This preset uses feedback rotation to rotate each consecutive echo 11.25° counter-clockwise in stero space (8 steps from right to left axis). The block matrix is used so that delay lines 1+2 and delay lines 3+4 are processed as identical stereo pairs.

Dispersing Echo: This preset uses geometry BOX_1 which has delay lengths not identical but very close together. The result is a temporal spread that increases recursively over time.

Laser Echo: This is the same principle as the previous preset but uses geometry DETUNED instead. The differently delayed echos are so close together they create a phasing effect rather than a dispersion.

Early Reflections: These presets use very short delay lengths with various geometies. The goal is to mimic *early reflections* as they typically happen from nearby walls and the floor. These early reflections additionally color and diffuse the sound, and are missing in the other reverb environments from above. To use early reflections, run one instance of AriesVerb doing early reflections feeding a second instance of AriesVerb doing the environment.

Special use

The presets ub this category show unusual things that can be done with the FDN algorithm.

E major −10c (Hymn): These preset has the delay lines tuned to an E-major chord off by 10 cents. Use this effect on the song "Hymn" by Barclay James Harvest. A slight modulation and a slight feedback rotation (3°|2°) are used to liven up the response.

LPF (change inv): This preset shows how a resonant low-pass filter can be constructed out of the FDN. This idea is to let the inverter control the resonant frequency by interpolating between even and odd resonance. The delay lengths are set to such a short length that the first odd resonant mode appears at 20 kHz. The rest of the setup exploits the fact that there must be a phase inversion across a resonant mode. Thus, if the dry signal is added to the wet signal, the spectrum above the resonant mode should cancel out. AriesVerb is such a great simulation of short delay lengths that the strategy actually works!

HPF (CHANGE INV): The same as the above with a different output vector. Now the spectrum below the resonant mode cancels out, giving a resonant high-pass filter.

AUTO WAH-WAH: A resonant low-pass filter with delay length modulation.

DC NOTCH (SHARP): The FDN has been turned into a notch filter. The wet signal is combined to the dry signal in such a way that the cancellation happens at the resonant peaks themselves, giving a sharp notch filter. In this case, one notch is at 0 Hz (DC), while the next is at 20 kHz.

DC NOTCH (FLAT): A play with the output vector, without any resonance. This is simply a subtraction of the wet signal from the dry signal, with a very short delay.

L/R → M/S: A play with the output vector, without any resonance and no delay. This preset converts left-right stereo to mid-side stereo. The output vector combines $1/\sqrt{2}(L+R)$ into the left channel and $1/\sqrt{2}(R-L)$ into the right channel.

M/S → L/R: Same as above, but the output vector now combines $1/\sqrt{2}(L-R)$ into the left channel and $1/\sqrt{2}(R+L)$ into the right channel.

INVERTER: This preset inverts the input signal.

INIT: This preset is a non-operation (identity).

Part IV
Credits

Built with GCC 4.3.1 on wxDev-C++ 6.10.2.
Documentation written on LyX 1.5.1 using MiKTeX 2.6.
VST Technology by Steinberg.

`http://www.ariescode.com`

License

References

[1] Schüler, Christian (2005) **"Fast Environmental Reverb Based On Feedback Delay Networks"**, in Game Programming Gems 5, Kim Pallister (editor), pp. 729 - 740, Charles River Media.

[2] Smith, Julius O. **"Physical Audio Signal Processing"**, available online at `http://ccrma.stanford.edu/~jos/pasp/pasp.html`.

[3] Jot, Jean Marc and Chaigne, Antoine (1991) **"Digital delay networks for designing artificial reverberators"**, in Proc. 90th Conv. Audio Eng. Soc., Feb. 1991, pre print 3030.