

Steve Harris' LADSPA Plugin Docs

steve@plugin.org.uk

2007-01-31

Contents

| | | |
|----------|--|----------|
| 1 | Preamble | 1 |
| 1.1 | What plugins? | 1 |
| 1.2 | Where can I get them | 1 |
| 2 | The plugins | 1 |
| 2.1 | Aliasing (alias, 1407) | 1 |
| 2.2 | Allpass delay line, noninterpolating (allpass_n, 1895) | 1 |
| 2.3 | Allpass delay line, linear interpolation (allpass_l, 1896) | 2 |
| 2.4 | Allpass delay line, cubic spline interpolation (allpass_c, 1897) | 2 |
| 2.5 | AM pitchshifter (amPitchshift, 1433) | 3 |
| 2.6 | Simple amplifier (amp, 1181) | 3 |
| 2.7 | Analogue Oscillator (analogueOsc, 1416) | 3 |
| 2.8 | Glame Bandpass Analog Filter (bandpass_a_iir, 1893) | 4 |
| 2.9 | Glame Bandpass Filter (bandpass_iir, 1892) | 4 |
| 2.10 | Bode frequency shifter (bodeShifter, 1431) | 4 |
| 2.11 | Bode frequency shifter (CV) (bodeShifterCV, 1432) | 5 |
| 2.12 | Glame Butterworth X-over Filter (bwxover_iir, 1902) | 5 |
| 2.13 | GLAME Butterworth Lowpass (butflow_iir, 1903) | 5 |
| 2.14 | GLAME Butterworth Highpass (butthigh_iir, 1904) | 5 |
| 2.15 | Chebyshev distortion (chebstortion, 1430) | 5 |
| 2.16 | Comb Filter (comb, 1190) | 5 |
| 2.17 | Comb delay line, noninterpolating (comb_n, 1889) | 5 |
| 2.18 | Comb delay line, linear interpolation (comb_l, 1887) | 6 |
| 2.19 | Comb delay line, cubic spline interpolation (comb_c, 1888) | 6 |
| 2.20 | Comb Splitter (combSplitter, 1411) | 7 |
| 2.21 | Constant Signal Generator (const, 1909) | 7 |
| 2.22 | Crossover distortion (crossoverDist, 1404) | 7 |
| 2.23 | DC Offset Remover (dcRemove, 1207) | 8 |
| 2.24 | Debug Plugin (debug, 1184) | 8 |
| 2.25 | Exponential signal decay (decay, 1886) | 8 |
| 2.26 | Decimator (decimator, 1202) | 8 |
| 2.27 | Declipper (declip, 1195) | 8 |
| 2.28 | Simple delay line, noninterpolating (delay_n, 1898) | 8 |
| 2.29 | Simple delay line, linear interpolation (delay_l, 1899) | 9 |

| | | |
|------|---|----|
| 2.30 | Simple delay line, cubic spline interpolation (delay_c, 1900) | 9 |
| 2.31 | Delayorama (delayorama, 1402) | 9 |
| 2.32 | Diode Processor (diode, 1185) | 10 |
| 2.33 | Audio Divider (Suboctave Generator) (divider, 1186) | 11 |
| 2.34 | DJ EQ (mono) (dj_eq_mono, 1907) | 11 |
| 2.35 | DJ EQ (dj_eq, 1901) | 11 |
| 2.36 | DJ flanger (djFlanger, 1438) | 12 |
| 2.37 | Dyson compressor (dysonCompress, 1403) | 12 |
| 2.38 | Fractionally Addressed Delay Line (fadDelay, 1192) | 13 |
| 2.39 | Fast Lookahead limiter (fastLookaheadLimiter, 1913) | 13 |
| 2.40 | Flanger (flanger, 1191) | 13 |
| 2.41 | FM Oscillator (fmOsc, 1415) | 14 |
| 2.42 | Foldover distortion (foldover, 1213) | 14 |
| 2.43 | Fast overdrive (foverdrive, 1196) | 15 |
| 2.44 | Frequency tracker (freqTracker, 1418) | 15 |
| 2.45 | Gate (gate, 1410) | 15 |
| 2.46 | Gate (gate, 1921) | 16 |
| 2.47 | Stereo Gate (stereo_gate, 1922) | 17 |
| 2.48 | Giant flange (giantFlange, 1437) | 18 |
| 2.49 | Gong model (gong, 1424) | 19 |
| 2.50 | Gong beater (gongBeater, 1439) | 22 |
| 2.51 | GSM simulator (gsm, 1215) | 22 |
| 2.52 | GVerb (gverb, 1216) | 23 |
| 2.53 | Hard Limiter (hardLimiter, 1413) | 24 |
| 2.54 | Harmonic generator (harmonicGen, 1220) | 24 |
| 2.55 | Hermes Filter (hermesFilter, 1200) | 25 |
| 2.56 | Glame Highpass Filter (highpass_iir, 1890) | 25 |
| 2.57 | Hilbert transformer (hilbert, 1440) | 25 |
| 2.58 | Impulse convolver (imp, 1199) | 25 |
| 2.59 | Nonbandlimited single-sample impulses (Frequency: Control) (impulse_fc, 1885) | 26 |
| 2.60 | Inverter (inv, 1429) | 27 |
| 2.61 | Karaoke (karaoke, 1409) | 27 |
| 2.62 | Artificial latency (artificialLatency, 1914) | 27 |
| 2.63 | L/C/R Delay (lcrDelay, 1436) | 27 |
| 2.64 | Lookahead limiter (lookaheadLimiter, 1435) | 28 |
| 2.65 | Lookahead limiter (fixed latency) (lookaheadLimiterConst, 1906) | 29 |
| 2.66 | Glame Lowpass Filter (lowpass_iir, 1891) | 29 |
| 2.67 | LS Filter (lsFilter, 1908) | 29 |
| 2.68 | Matrix: MS to Stereo (matrixMSSt, 1421) | 30 |
| 2.69 | Matrix Spatialiser (matrixSpatialiser, 1422) | 30 |
| 2.70 | Matrix: Stereo to MS (matrixStMS, 1420) | 30 |
| 2.71 | Multiband EQ (mbeq, 1197) | 30 |
| 2.72 | Modulatable delay (modDelay, 1419) | 30 |
| 2.73 | Multivoice Chorus (multivoiceChorus, 1201) | 31 |
| 2.74 | Mag's Notch Filter (notch_iir, 1894) | 31 |
| 2.75 | LFO Phaser (lfoPhaser, 1217) | 31 |

| | | |
|-------|--|----|
| 2.76 | 4 x 4 pole allpass (fourByFourPole, 1218) | 31 |
| 2.77 | Auto phaser (autoPhaser, 1219) | 31 |
| 2.78 | Pitch Scaler (pitchScale, 1193) | 31 |
| 2.79 | Higher Quality Pitch Scaler (pitchScaleHQ, 1194) | 32 |
| 2.80 | Plate reverb (plate, 1423) | 32 |
| 2.81 | Pointer cast distortion (pointerCastDistortion, 1910) | 32 |
| 2.82 | Rate shifter (rateShifter, 1417) | 33 |
| 2.83 | Retro Flanger (retroFlange, 1208) | 33 |
| 2.84 | Reverse Delay (5s max) (revdelay, 1605) | 33 |
| 2.85 | Ringmod with two inputs (ringmod_2i1o, 1188) | 34 |
| 2.86 | Ringmod with LFO (ringmod_1i1o1l, 1189) | 34 |
| 2.87 | Barry's Satan Maximiser (satanMaximiser, 1408) | 34 |
| 2.88 | SC1 (sc1, 1425) | 35 |
| 2.89 | SC2 (sc2, 1426) | 35 |
| 2.90 | SC3 (sc3, 1427) | 36 |
| 2.91 | SC4 (sc4, 1882) | 37 |
| 2.92 | SC4 mono (sc4m, 1916) | 38 |
| 2.93 | SE4 (se4, 1883) | 39 |
| 2.94 | Wave shaper (shaper, 1187) | 40 |
| 2.95 | Signal sifter (sifter, 1210) | 40 |
| 2.96 | Sine + cosine oscillator (sinCos, 1881) | 40 |
| 2.97 | Single band parametric (singlePara, 1203) | 40 |
| 2.98 | Sinus wavewrapper (sinusWavewrapper, 1198) | 41 |
| 2.99 | Smooth Decimator (smoothDecimate, 1414) | 41 |
| 2.100 | Mono to Stereo splitter (split, 1406) | 41 |
| 2.101 | Step Demuxer (stepMuxer, 1212) | 41 |
| 2.102 | Surround matrix encoder (surroundEncoder, 1401) | 42 |
| 2.103 | State Variable Filter (svf, 1214) | 43 |
| 2.104 | Tape Delay Simulation (tapeDelay, 1211) | 44 |
| 2.105 | Transient mangler (transient, 1206) | 44 |
| 2.106 | Triple band parametric with shelves (triplePara, 1204) | 44 |
| 2.107 | Valve saturation (valve, 1209) | 44 |
| 2.108 | Valve rectifier (valveRect, 1405) | 44 |
| 2.109 | VyNil (Vinyl Effect) (vynil, 1905) | 44 |
| 2.110 | Wave Terrain Oscillator (waveTerrain, 1412) | 45 |
| 2.111 | Crossfade (xfade, 1915) | 45 |
| 2.112 | Crossfade (4 outs) (xfade4, 1917) | 45 |
| 2.113 | z-1 (zm1, 1428) | 45 |

3 Licensing 46

1 Preamble

1.1 What plugins?

This is the documentation for some plugins that I have written for the *Linux Audio Developers Simple Plugin Architecture*. It is a nice audio plugin architecture with a very easy learning curve.

1.2 Where can I get them

From the website at <http://plugin.org.uk/>.

2 The plugins

2.1 Aliasing (alias, 1407)

Simulates aliasing using nyquist frequency modulation. Produces wacky results if the blocks aren't even numbers of samples long.

Aliasing level

Controls the amount of simulated aliasing in the output.

2.2 Allpass delay line, noninterpolating (allpass_n, 1895)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.3 Allpass delay line, linear interpolation (allpass_l, 1896)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.4 Allpass delay line, cubic spline interpolation (allpass_c, 1897)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.5 AM pitchshifter (amPitchshift, 1433)

This plugin works by running a single write pointer (monotonic) and two read pointers (pitchscaled) over a ringbuffer. The output is faded between the two read pointers according to the sine of the distance from the write pointer. The design is based on the mechanism of a mechanical pitchshifter I saw in the Gemeentemuseum in Den Haag, though I'm sure it is a common enough algorithm.

Pitch shift

The multiple of the output pitch, eg. 2.0 will increase the pitch by one octave.

Buffer size

The order of magnitude of the buffer size. Small buffers will sound fluttery, large buffers will have flangy sounding echoes. I recommend a buffer size of 3 for a reasonable compromise, with wideband material at around 48KHz. For drums you might have to lower it, and for voiced background noises it can go higher.

2.6 Simple amplifier (amp, 1181)

Amps gain (dB)

Controls the gain of the input signal in dB's.

2.7 Analogue Oscillator (analogueOsc, 1416)

This plugin simulates the output you get from an analogue synth's oscillators. You can get a reasonable emulation of a 303's square (for example) if you set the warmth to about 0.4 and the instability to about 0.05. The frequency is currently a control input, and there is no interpolation, so if your host is using large block size it will sound stepy. I'm unsure whether to convert this to an audio input or interpolate the control in.

Waveform (1=sin, 2=tri, 3=squ, 4=saw)

The basic shape of the waveform is selected using this control:

| Value | Waveform |
|-------|----------|
| 1 | Sine |
| 2 | Triangle |
| 3 | Square |
| 4 | Saw |

Frequency (Hz)

The frequency of the output (Hz).

Warmth

The degree of softening that is applied to the generated waveform, reduces the number of harmonics in the output.

Instability

The degree of pitch instability of the output. Turning this too high with square and saw waves will produce an annoying jittery sound, I want to fix this but it is tricky.

2.8 Glame Bandpass Analog Filter (bandpass_a_iir, 1893)

IIR bandpass filter modeled after an analog circuit. This filter was ported from the glame multitrack editor to ladspa.

2.9 Glame Bandpass Filter (bandpass_iir, 1892)

IIR bandpass filter based using chebishev coefficients. The filter allows you to tweak the number of stages used for filtering. Every stage adds two more poles, which leads to a steeper dropoff. More stages need more CPU power. This filter was ported from the glame multitrack editor to ladspa.

2.10 Bode frequency shifter (bodeShifter, 1431)

A Bode/Moog Frequency Shifter is a popular analogue synth module, it works by shifting all the frequencies of an input signal up or down by a specified frequency. This version shifts in both directions at the same time as it almost no extra work and it's often useful to have both directions. It doesn't actually work in the same way as an analogue Bode/Moog, which use Dome filters as the core, it uses a Hilbert Transformer, which is much simpler to implement in digital systems. The output is very similar though, and people are familiar with the name Bode. The theory of operation is pretty simple, and uses some clever maths to cancel out the upper or lower sidebands of a ringmodulator applied to the input signal. Read the source if you want more information. The Hilbert Transformer coefficients came from mkfilter, the excellent filter calculator, available at <http://www-users.cs.york.ac.uk/~fisher/mkfilter/>.

Frequency shift

Controls the frequency shift applied to the input signal, in Hz. Note, this is not a pitch shift, so you not get natural sounding results out, it is an audio effect popular with modular synthesists.

2.11 Bode frequency shifter (CV) (bodeShifterCV, 1432)

See the non CV version for information. This is more or less a copy of the Doepfer A126, <http://www.doepfer.de/a126.htm>.

Shift CV

Controls the frequency shift applied to the input signal, in KHz.

2.12 Glame Butterworth X-over Filter (bwxover_iir, 1902)

Butterworth X-over filter

2.13 GLAME Butterworth Lowpass (buttlow_iir, 1903)

Butterworth lowpass filter

2.14 GLAME Butterworth Highpass (butthigh_iir, 1904)

Butterworth highpass filter

2.15 Chebyshev distortion (chebstortion, 1430)

This is an interesting distortion effect that is seeded from incoming signal envelope. As the level of the signal rises more and more harmonics will be added to the output signal. The distortion control sets the sensitivity of the input. The effect evolved from some experiments between Tim Goetze and myself, attempting to emulate valve based guitar amp distortion. This was one of the failures, but it still makes an interesting noise.

2.16 Comb Filter (comb, 1190)

Band separation (Hz)

Controls the distance between the filter peaks.

Feedback

Feedback level, increases the distinctive wooshy phaser sound.

2.17 Comb delay line, noninterpolating (comb_n, 1889)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.18 Comb delay line, linear interpolation (comb_l, 1887)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.19 Comb delay line, cubic spline interpolation (comb_c, 1888)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

Decay Time (s)

Time for the echoes to decay by 60 decibels. If this time is negative then the feedback coefficient will be negative, thus emphasizing only odd harmonics at an octave lower.

2.20 Comb Splitter (combSplitter, 1411)

Divides the input up into two parts with frequency peaks at f Hz intervals, skewed by $f/2$ Hz between the two outputs. Mixing the two outputs will get you exactly the input signal. I generally use this trick to divide up an input signal, process the two halves differently, then mix them again. It sounds pretty funky.

Band separation (Hz)

The distance between the frequency peaks.

Output 1

The sum output.

Output 2

The difference output.

2.21 Constant Signal Generator (const, 1909)

This plugin add an output DC offset at the given amplitude to the input signal. It has no real use other than for debugging and in modular synths.

Signal amplitude

Controls the amplitude of the output signal.

2.22 Crossover distortion (crossoverDist, 1404)

This is a simulation of the distortion that happens in class B and AB power amps when the signal crosses 0. For class B simulations the smooth value should be set to about 0.3 +/- 0.2 and for AB it should be set to near 1.0.

Crossover amplitude

Controls the point at which the output signal becomes linear.

Smoothing

Controls degree of smoothing of the crossover point.

2.23 DC Offset Remover (dcRemove, 1207)

Simply removes the DC (0 Hz) component from an audio signal, uses a high pass filter, so has some side effects, but they should be minimal.

2.24 Debug Plugin (debug, 1184)

Prints some stats about the input stream to stdout. Not intended for general use.

2.25 Exponential signal decay (decay, 1886)

Based on work by James McCartney in SuperCollider.

Decay Time (s)

Time for the echoes to decay by 60 decibels.

2.26 Decimator (decimator, 1202)

Decimates (reduces the effective sample rate), and reduces the bit depth of the input signal, allows non integer values for smooth transitions between clean and lofi signals.

Bit depth

The bit depth that the signal will be reduced to.

Sample rate (Hz)

The sample rate that the signal will be resampled at.

2.27 Declipper (declip, 1195)

Removes nasty clicks from input signals, not very kind to them though. This code came from the music-dsp mailing list, but it was unattributed, if it's yours, please drop me a line and I'll credit you.

2.28 Simple delay line, noninterpolating (delay_n, 1898)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

2.29 Simple delay line, linear interpolation (delay_l, 1899)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

2.30 Simple delay line, cubic spline interpolation (delay_c, 1900)

Based on work by James McCartney in SuperCollider.

Max Delay (s)

Maximum delay. Used to set the delay buffer size upon activation. Cannot be modulated. Note that if you do not connect to this port before activation, it will default to 1 second.

2.31 Delayorama (delayorama, 1402)**Random seed**

Controls the random numbers that will be used to stagger the delays and amplitudes if random is turned up on them. Changing this forces the random values to be recalculated.

Input gain (dB)

Controls the gain of the input signal in dB's.

Feedback (%)

Controls the amount of output signal fed back into the input.

Number of taps

Controls the number of taps in the delay.

First delay (s)

The time of the first delay.

Delay range (s)

The time difference between the first and last delay.

Delay change

The scaling factor between one delay and the next.

Delay random (%)

The random factor applied to the delay.

Amplitude change

The scaling factor between one amplitude and the next.

Amplitude random (%)

The random factor applied to the amplitude.

Dry/wet mix

The level of delayed sound mixed into the output.

2.32 Diode Processor (diode, 1185)

Mangles the signal as if it had been passed through a diode rectifier network. You should probably follow this with a DC offset remover, unless you want the offset.

Mode (0 for none, 1 for half wave, 2 for full wave)

The mode parameter is continuously variable from thru to half-wave rectification to full-wave to silence.

2.33 Audio Divider (Suboctave Generator) (divider, 1186)

Reduces the period of the signal by the factor given, and makes it a square wave in the process. Has some amplitude tracking capability, but not really useful on complex signals.

Denominator

The factor the incoming frequency will be divided by.

2.34 DJ EQ (mono) (dj_eq_mono, 1907)

The design for this plugin is taken from the Allen & Heath Xone 32 DJ mixer. It was suggested by Patrick Shirkey. Mono version requested by Adam King

Lo gain (dB)

Controls the gain of the low (100Hz) peak/dip band

Mid gain (dB)

Controls the gain of the mid (1000Hz) peak/dip band

Hi gain (dB)

Controls the gain of the high (10000Hz) shelf band

2.35 DJ EQ (dj_eq, 1901)

The design for this plugin is taken from the Allen & Heath Xone 32 DJ mixer. It was suggested by Patrick Shirkey.

Lo gain (dB)

Controls the gain of the low (100Hz) peak/dip band

Mid gain (dB)

Controls the gain of the mid (1000Hz) peak/dip band

Hi gain (dB)

Controls the gain of the high (10000Hz) shelf band

2.36 DJ flanger (djFlanger, 1438)

This is a flanger which is more or less typical of DJ mixing desks. Requested by Patrick Shirkey.

LFO sync

When turned from off to on it resets the phase of the LFO back to the start of the cycle. Used to sync the LFO to the track.

LFO period (s)

The cycle period of the LFO in seconds.

LFO depth (ms)

The maximum delay the LFO will use to flange, in milliseconds.

Feedback (%)

The amount of the delays output that is mixed back into the delay.

2.37 Dyson compressor (dysonCompress, 1403)**Peak limit (dB)**

Controls the desired limit of the output signal in dB's.

Release time (s)

Controls the time taken for the compressor to relax its gain control over the input signal.

Fast compression ratio

I have no clear idea what this controls.

Compression ratio

I have no clear idea what this controls.

2.38 Fractionally Addressed Delay Line (fadDelay, 1192)

A fixed ring buffer delay implementation. Has different dynamics to a normal delay, more suitable for certain things. Changes in delay length are generally more pleasing, but delays >2s long have reduced sound quality.

Delay (seconds)

The neutral delay time is 2 seconds. Times above 2 seconds will have reduced quality and times below will have increased CPU usage.

2.39 Fast Lookahead limiter (fastLookaheadLimiter, 1913)

This is a limiter with an attack time of 5ms. It adds just over 5ms of latency to the input signal, but it guarantees that there will be no signals over the limit, and tries to get the minimum amount of distortion.

Input gain (dB)

Gain that is applied to the input stage. Can be used to trim gain to bring it roughly under the limit or to push the signal against the limit.

Limit (dB)

The maximum output amplitude. Peaks over this level will be attenuated as smoothly as possible to bring them as close as possible to this level.

Release time (s)

The time taken for the limiters attenuation to return to 0 dB's

Attenuation (dB)

The current attenuation of the signal coming out of the delay buffer.

2.40 Flanger (flanger, 1191)

A digital flanger implementation. Uses a novel zero excursion, controlled bandwidth modulation function, which should make the modulation less repetitive and noticeable. This effect is similar in character to a phaser (see section 2.75). The main difference is that a phaser sounds more regular and stable.

Delay base (ms)

This is the offset from the input time that the detune delay moves around. 10 is probably a good starting value.

Max slowdown (ms)

This is the maximum delay that will be applied to the delayed signal, relative to the dry signal.

LFO frequency (Hz)

This is the core frequency that the 'LFO' will move at. The LFO isn't actually an oscillator, but it does vary periodically.

Feedback

Feedback applied from the output to the input, increases the depth of the effect, but makes it sound less like a real flanger.

2.41 FM Oscillator (fmOsc, 1415)

Waveform (1=sin, 2=tri, 3=squ, 4=saw)

The shape of the waveform.

| Value | Waveform |
|-------|----------|
| 1 | Sine |
| 2 | Triangle |
| 3 | Square |
| 4 | Saw |

Frequency (Hz)

The frequency of the output (in Hertz).

2.42 Foldover distortion (foldover, 1213)

Uses a sinwave approximation to simulate valve style foldover distortion. Probably should have a DC offset remover on the output, but it's not always necessary.

Drive

Controls the degree of distortion.

Skew

Controls the asymmetry of the waveform.

2.43 Fast overdrive (foverdrive, 1196)

A simple overdrive. Compresses the extreme peaks to make a sound similar to an overdriven amplifier.

Drive level

Controls the point at which the signal starts to distort, and the degree of distortion.

2.44 Frequency tracker (freqTracker, 1418)

Tracking speed

This controls the level of damping applied to the output. High values will make the frequency output jump around, low values will make it a bit slow to respond.

2.45 Gate (gate, 1410)

The parameters are copied from the Drawmer DS-201, but I've never used one, so if someone out there has one, please tell me if it behaves differently.

LF key filter (Hz)

Controls the cutoff of the low frequency filter (highpass).

HF key filter (Hz)

Controls the cutoff of the high frequency filter (lowpass).

Threshold (dB)

Controls the level at which the gate will open.

Attack (ms)

Controls the time the gate will take to open fully.

Hold (ms)

Controls the minimum time the gate will stay open for.

Decay (ms)

Controls the time the gate will take to close fully.

Range (dB)

Controls the difference between the gate's open and closed state.

Output select (-1 = key listen, 0 = gate, 1 = bypass)

Controls output monitor. -1 is the output of the key filters (so you can check what is being gated on). 0 is the normal, gated output. 1 is bypass mode.

2.46 Gate (gate, 1921)

The parameters are copied from the Drawmer DS-201, but I've never used one, so if someone out there has one, please tell me if it behaves differently.

LF key filter (Hz)

Controls the cutoff of the low frequency filter (highpass).

HF key filter (Hz)

Controls the cutoff of the high frequency filter (lowpass).

Key level (dB)

Shows the current level of the key.

Threshold (dB)

Controls the level at which the gate will open.

Attack (ms)

Controls the time the gate will take to open fully.

Hold (ms)

Controls the minimum time the gate will stay open for.

Decay (ms)

Controls the time the gate will take to close fully.

Range (dB)

Controls the difference between the gate's open and closed state.

Output select (-1 = key listen, 0 = gate, 1 = bypass)

Controls output monitor. -1 is the output of the key filters (so you can check what is being gated on). 0 is the normal, gated output. 1 is bypass mode.

2.47 Stereo Gate (stereo_gate, 1922)

Derived from Steve Harris' gate plugin

LF key filter (Hz)

Controls the cutoff of the low frequency filter (highpass).

HF key filter (Hz)

Controls the cutoff of the high frequency filter (lowpass).

Key level (dB)

Shows the current level of the key.

Threshold (dB)

Controls the level at which the gate will open.

Attack (ms)

Controls the time the gate will take to open fully.

Hold (ms)

Controls the minimum time the gate will stay open for.

Decay (ms)

Controls the time the gate will take to close fully.

Range (dB)

Controls the difference between the gate's open and closed state.

Output select (-1 = key listen, 0 = gate, 1 = bypass)

Controls output monitor. -1 is the output of the key filters (so you can check what is being gated on). 0 is the normal, gated output. 1 is bypass mode.

2.48 Giant flange (giantFlange, 1437)

This is a fairly normal flanger but with excessively long delay times. Requested by Patrick Shirkey. To cut down the memory requirements the internal delay buffer only has 15bits of resolution, so there is no headroom, if you feed in signals over 0dB it will clip the output. There is code to soften the effect of the clipping, but beware of it.

Double delay

doubles the length of the delays, this will reduce the sound quality.

LFO frequency 1 (Hz)

The delay of the first LFO in seconds.

Delay 1 range (s)

The delay range of the first LFO in seconds.

LFO frequency 2 (Hz)

The delay of the second LFO in seconds.

Delay 2 range (s)

The delay range of the second LFO in seconds.

Feedback

The amount of the delays output that is mixed back into the delay.

Dry/Wet level

The amounts of the input and effect mixed to produce the output.

2.49 Gong model (gong, 1424)

A physical model of a metal gong. Based on Josep Comajuncosas' gong explorer, which was built in Sync Modular, it uses 8 linear waveguides with nonlinear filters to model the gong surface.

Inner damping

Controls the degree of damping in the centre of the gong.

Outer damping

Controls the degree of damping on the edge of the gong.

Mic position

Controls the vertical position of the "microphone", 0 is the centre and 1 is the edge.

Inner size 1

The size of the upper, inner waveguide.

Inner stiffness 1 +

The stiffness of the gong against deflections in the positive direction.

Inner stiffness 1 -

The stiffness of the gong against deflections in the negative direction.

Inner size 2

The size of the right, inner waveguide.

Inner stiffness 2 +

The stiffness of the gong against deflections in the positive direction.

Inner stiffness 2 -

The stiffness of the gong against deflections in the negative direction.

Inner size 3

The size of the lower, inner waveguide.

Inner stiffness 3 +

The stiffness of the gong against deflections in the positive direction.

Inner stiffness 3 -

The stiffness of the gong against deflections in the negative direction.

Inner size 4

The size of the left, inner waveguide.

Inner stiffness 4 +

The stiffness of the gong against deflections in the positive direction.

Inner stiffness 4 -

The stiffness of the gong against deflections in the negative direction.

Outer size 1

The size of the upper right, outer waveguide.

Outer stiffness 1 +

The stiffness of the gong against deflections in the positive direction.

Outer stiffness 1 -

The stiffness of the gong against deflections in the negative direction.

Outer size 2

The size of the lower right, outer waveguide.

Outer stiffness 2 +

The stiffness of the gong against deflections in the positive direction.

Outer stiffness 2 -

The stiffness of the gong against deflections in the negative direction.

Outer size 3

The size of the lower left, outer waveguide.

Outer stiffness 3 +

The stiffness of the gong against deflections in the positive direction.

Outer stiffness 3 -

The stiffness of the gong against deflections in the negative direction.

Outer size 4

The size of the upper left, outer waveguide.

Outer stiffness 4 +

The stiffness of the gong against deflections in the positive direction.

Outer stiffness 4 -

The stiffness of the gong against deflections in the negative direction.

2.50 Gong beater (gongBeater, 1439)

A plugin to simulator the action of a beater on a gong surface, used to trigger the gong physical model. It is triggered by an impulse on the input, eg. from a mic or piezo placed near a solid surface, for an event sequencer.

Impulse gain (dB)

The gain of the input impulse mixed into the output, bringing this up allows you to make the outputted strike more impulsive, but may reduce the gongyness of the resulting output sound.

Strike gain (dB)

The gain of the simulated pressure wave mixed into the output, bringing this up allows you to make the outputted strike more pure. The final output level is also proportional to the amplitude of the trigger.

Strike duration (s)

The duration of the pressure wave used to simulate the action of the beater on the gong surface. The longer the duration the more sonorous the resulting gong sound.

2.51 GSM simulator (gsm, 1215)

Encodes and decodes a signal using the GSM voice compression system. Has the effect of making the signal sound like it is being sent over a European mobile phone network.

Dry/wet mix

Controls the dry/wet mix, 0 will give you the dry signal (but with the appropriate amount of delay), 1 will give you a totally wet signal.

Number of passes

The number of times the signal is sent through the encode/decode process. Increases the CPU consumption almost linearly, and it will become more peaky so less friendly to realtime processing.

Error rate (bits/block)

The number of simulated bits that get changed during the transmission process. I really wanted to reduce the bandwidth to get that "shouting down a drainpipe" effect, but I'm not sure how the reduced bandwidth is dealt with by real phones. I suspect it's heavily patented technology.

2.52 GVerb (gverb, 1216)

A mono in, stereo out reverb implementation by Juhana Sadeharju (kouhia at nic.funet.fi). I ported it to LADSPA and did some testing. Please contact Juhana directly regarding any bugs you find. Paul Winkler recommends a good starting point for a "large hall" as follows:

| Parameter | amp; Value |
|------------------------|------------|
| Roomsize | amp; 200 |
| Reverb time | amp; 1.3 |
| Damping | amp; 0.4 |
| Input bandwidth | amp; 0.5 |
| Dry signal level | amp; 0 |
| Early reflection level | amp; -12 |
| Tail level | amp; 0 |

Roomsize (m)

The size of the room, in meters. Excessivly large, and excessivly small values will make it sound a bit unrealistic. Values of around 30 sound good.

Reverb time (s)

Reverb decay time, in seconds. 7 is a good place to start.

Damping

This controls the high frequency damping (a lowpass filter), values near 1 will make it sound very bright, values near 0 will make it sound very dark.

Input bandwidth

This is like a damping control for the input, it has a similar effect to the damping control, but is subtly different.

Dry signal level (dB)

The amount of dry signal to be mixed with the reverberated signal.

Early reflection level (dB)

The quantity of early reflections (scatter reflections directly from the source). Think of Lexicons ambiance patches.

Tail level (dB)

The level of the classic reverb tail reflections.

2.53 Hard Limiter (hardLimiter, 1413)

Brick hard limiter with residue mixer.

Wet level

Output level for limited signal.

Residue level

Output level for residue signal.

2.54 Harmonic generator (harmonicGen, 1220)**What does it do?**

Allows you to add harmonics and remove the fundamental from any audio signal.

Known bugs

There is no bandwidth limiting filter on the output, so it is easy to create excessively high frequency harmonics that could cause aliasing problems. In practice this doesn't seem to be a serious problem however.

Examples

There are many interesting effects you can achieve with sinewaves, one example is producing bandlimited squarewaves from sinewaves. To do this set the parameters to 1, 0, -0.3333, 0, 0.2, 0, -0.14285, 0, 0.11111. To get a triangle like signal use 1, 0, -0.3333, 0, -0.2, 0, -0.14285, 0, -0.11111.

Fundamental magnitude

The amplitude of the fundamental of the signal, reduce it to 0 to remove the base signal altogether, or -1 to phase invert it.

2nd harmonic magnitude

The 2nd harmonic, its frequency is twice the frequency of the harmonic. Even harmonics add a distorted feel to the sound, valve (tube) amplifiers introduce distortions at all the harmonics.

3rd harmonic magnitude

The 3rd harmonic, its frequency is three times the frequency of the fundamental. Transistor amplifiers only introduce distortion into the odd harmonics.

2.55 Hermes Filter (hermesFilter, 1200)

This plugin is a simulation of a modern analogue synth called a Pro Tone, with some extra features bolted on, like a crossover. I tried to make it as comprehensive as possible, without requiring ludicrous amounts of CPU juice. N.B. as far as I know, no one has tried to use this (I certainly haven't), so it may be full of bugs and what not. The parameters are all undocumented, but there is a diagram of the routing on the website. Without a custom interface however it would be very hard to use. Historical note: the name is a bad pun, it comes from the name Hermes Trimegistus given to the Egyptian god Thoth by the Greeks, it means Thrice Blessed, or something similar.

2.56 Glame Highpass Filter (highpass_iir, 1890)

IIR highpass filter based using chebishev coefficients. The filter allows you to tweak the number of stages used for filtering. Every stage adds two more poles, which leads to a steeper dropoff. More stages need more CPU power. This filter was ported from the glame multitrack editor to ladspa.

2.57 Hilbert transformer (hilbert, 1440)

A Hilbert Transformer phase shifts the input signal by 90degrees. It outputs the 90 degree phase shifted signal and the unshifted signal, both delayed by an equivalent amount. This plugin was written for a demo at the LAD Meet in 2003.

2.58 Impulse convolver (imp, 1199)

This is a convolver for a set of fairly short impulses.

The set of impulses has to be compiled in, they are:

| Id | Impulse |
|----|--|
| 1 | Unit impulse (identity) |
| 2 | My flat (light natural reverb) |
| 3 | Yamaha Marshall stack simulator |
| 4 | Fender 68 Vibrolux (SM57 on axis) |
| 5 | Fender 68 Vibrolux (SM57 off axis) |
| 6 | Fender 68 Vibrolux (Audio-technica AT4050) |
| 7 | Fender 68 Vibrolux (Neumann U87) |
| 8 | Fender Bassman (SM57 on axis) |
| 9 | Fender Bassman (SM57 off axis) |
| 10 | Fender Bassman (Audio-technica AT4050) |
| 11 | Fender Bassman (Neumann U87) |
| 12 | Fender Superchamp (SM57 on axis) |
| 13 | Fender Superchamp (SM57 off axis) |
| 14 | Fender Superchamp (Audio-technica AT4050) |
| 15 | Fender Superchamp (Neumann U87) |
| 16 | Marshall JCM2000 (SM57 on axis) |
| 17 | Marshall JCM2000 (SM57 off axis) |
| 18 | Marshall Plexi (SM57 on axis) |
| 19 | Marshall Plexi (SM57 off axis) |
| 20 | Matchless Chieftan (SM57 on axis) |
| 21 | Matchless Chieftan (SM57 off axis) |

The first three were quickly grabbed by me using jack_impulse_grabber, and the others we collected by someone else, but unfortunately I've lost his email address and can't find him on the web :(

Impulse ID

Selects the impulse to convolve with. New impulses have to be compiled in.

High latency mode

If you are running with blocks that are not whole powers of two long, or you are hearing distortion, try changing this to 1.

Gain (dB)

Controls the gain of the output signal in dB's.

2.59 Nonbandlimited single-sample impulses (Frequency: Control) (impulse_fc, 1885)

Based on work by James McCartney in SuperCollider.

Frequency (Hz)

Frequency for the impulses.

2.60 Inverter (inv, 1429)

A utility plugin that inverts the signal, also (wrongly) known as a 180 degree phase shift.

2.61 Karaoke (karaoke, 1409)

Attempts to strip the vocals from a stereo signal.

Vocal volume (dB)

Controls the attenuation of the vocal (centre channel) in dB's. The greater the attenuation the greater the loss of stereo field.

2.62 Artificial latency (artificialLatency, 1914)

Reports its delay value as systemic latency. Does nothing else, *this is not a delay*. Can be used to correct for latency between channels.

2.63 L/C/R Delay (lcrDelay, 1436)

This is a left/centre/right delay with feedback, based on the one in the Korg Trinity.
Requested by Marek Peteraj.

L delay (ms)

The delay of the left output in milliseconds.

L level

The level of the left output.

C delay (ms)

The delay of the centre output in milliseconds.

C level

The level of the centre output.

R delay (ms)

The delay of the right output in milliseconds.

R level

The level of the right output.

Feedback

The amount of the C delay output that is mixed back into the delay.

High damp (%)

The damping of the high frequencies in the feedback path.

Low damp (%)

The damping of the low frequencies in the feedback path.

Spread

The width of the stereo image.

Dry/Wet level

The amounts of the input and effect mixed to produce the output.

2.64 Lookahead limiter (lookaheadLimiter, 1435)

Limit (dB)

The maximum output amplitude. Peaks over this level will be attenuated as smoothly as possible to bring them as close as possible to this level.

Lookahead delay

The delay used by the lookahead predictor. The longer the delay the smoother the limiting will be, but higher the latency.

Attenuation (dB)

The current attenuation of the signal coming out of the delay buffer.

2.65 Lookahead limiter (fixed latency) (lookaheadLimiterConst, 1906)

A lookahead limiter - similar to the original Lookahead Limiter, but with a constant latency of around 150ms and a reduced maximum lookahead time.

Limit (dB)

The maximum output amplitude. Peaks over this level will be attenuated as smoothly as possible to bring them as close as possible to this level.

Lookahead time (s)

The lookahead time used by the lookahead predictor. The longer the time the smoother the limiting will be, but will tend to make the changes in dynamic range more obvious.

Attenuation (dB)

The current limiting attenuation of the signal coming out of the delay buffer.

2.66 Glame Lowpass Filter (lowpass_iir, 1891)

IIR lowpass filter based using chebishev coefficients. The filter allows you to tweak the number of stages used for filtering. Every stage adds two more poles, which leads to a steeper dropoff. More stages need more CPU power. This filter was ported from the glame multitrack editor to ladspa.

2.67 LS Filter (lsFilter, 1908)

This is a filter created for the LinkSampler project - its designed to closely follow the filter used in giga sampler.

Filter type (0=LP, 1=BP, 2=HP)

The type of the filter, 0 for low pass, 1 for band pass, 2 for high pass.

Cutoff frequency (Hz)

Controls the frequency at which the filter starts to effect the audio signal.eg. a low-pass filter with a cutoff frequency of 1000 Hz will only let frequencies below 1000 Hz through.

Resonance

Creates a peak at the cutoff frequency, for the classic overdriven filter sound. At high values the peak at the cutoff will overwhelm the filtered signal.

2.68 Matrix: MS to Stereo (matrixMSSt, 1421)

Width

The width of the dematrixed stereo field. 1 will give you normal width, 0 will make it completely mono, < 1 will make it narrower and > 1 will make it wider.

2.69 Matrix Spatialiser (matrixSpatialiser, 1422)

A simple spatializer that lets you control the width of a stereo signal.

We convert it into a MS (mid/side) signal, manipulate the gain coefficients with a constant-power panning function, and reconvert to left/right stereo.

$$mid = (i_left + i_right)/2$$

$$side = (i_left - i_right)/2$$

$$width = (-\pi/4)..0..(\pi/4)$$

$$o_left = mid * \cos(width + \pi/4) \quad o_right = side * \sin(width + \pi/4)$$

shifted by $\pi/4$, so that 0 is neutral.

Width

0 is neutral (unmodified signal) + 128 is side only (=very wide) - 128 is mid only (=mono)

2.70 Matrix: Stereo to MS (matrixStMS, 1420)

2.71 Multiband EQ (mbeq, 1197)

This is a fairly typical multiband graphical equalizer. It's implemented using a FFT, so it takes quite a lot of CPU power, but should have less phase effects than an equivalent filter implementation. If the input signal is at too low a sample rate then the top bands will be ignored, the highest useful band will always be a high shelf.

2.72 Modulatable delay (modDelay, 1419)

A delay whose tap can be modulated at audio rate. Requested by Matthias Nagorni at LinuxTag 2002, in order to make a Leslie simulator.

2.73 Multivoice Chorus (multivoiceChorus, 1201)

This is an implementation of a Multivoice (as opposed to Multiscale) chorus algorithm. It uses a novel, sinc based noise interpolation method to produce a subtle modulation law which makes it possible to get away with larger numbers of voices without the metallic, artificial sound common in chorus effects.

Voice separation (ms)

The individual voices can either be running at the same base delay (set this to zero) or staggered. Setting this to non-zero values can make the output sound richer, but will make it sound grainy with some type of signal.

Detune (%)

The maximum amount that a voice will be detuned by. I recommend a value of 1, but you may be able to get away with higher values if the signal is less harmonic.

LFO frequency (Hz)

The frequency that the detune effect will be modulated at. A matter of taste, for most types of input lower will be more subtle.

Output attenuation (dB)

With large numbers of voices the output can become too high, so use this to trim the amplitude to a more helpful level.

2.74 Mag's Notch Filter (notch_iir, 1894)

IIR notch filter based using chebishev coefficients. The filter allows you to tweak the number of stages used for filtering. Every stage adds two more poles, which leads to a steeper dropoff. More stages need more CPU power.

2.75 LFO Phaser (lfoPhaser, 1217)

2.76 4 x 4 pole allpass (fourByFourPole, 1218)

2.77 Auto phaser (autoPhaser, 1219)

2.78 Pitch Scaler (pitchScale, 1193)

A pitch shifter implementation that scales the harmonics appropriately with the base frequencies. It is an implementation of Stephen M. Sprengler's pitch scaler design. It gives reasonable, general purpose results for small changes, but won't give Antares or Eventide anything to worry about. The FFT block size and oversampling has been kept at a minimum to keep the CPU usage low.

2.79 Higher Quality Pitch Scaler (pitchScaleHQ, 1194)

A pitch shifter implementation that scales the harmonics appropriately with the base frequencies. It is an implementation of Stephen M. Sprengler's pitch scaler design. It gives reasonable, general purpose results for small changes, but won't give Antares or Eventide anything to worry about. The FFT block size and oversampling has been kept at reasonable levels to keep the CPU usage low, but it is smoother than the other Pitch Scaler.

Pitch co-efficient

The pitch scaling factor, a value of 2.0 will increase the pitch by one octave, etc.

2.80 Plate reverb (plate, 1423)

A physical model of a steel plate reverb. Based on Josep Comajuncosas' gong model, it uses 8 linear waveguides to model the plate.

Reverb time

Controls the RT60 time of the reverb. Actually controls the size of the plate. The mapping between plate size and RT60 time is just a heuristic, so it's not very accurate.

Damping

Controls the degree that the surface of the plate is damped.

Dry/wet mix

Controls the balance between the dry and wet signals.

2.81 Pointer cast distortion (pointerCastDistortion, 1910)

This distortion is created by treating the floating point representation of the input signal as a 0.32 1's complement fixedpoint integer. Its very unmusical but surprisingly recognisable. I'm not sure that its useful for anything, but it can make interesting noises.

Effect cutoff freq (Hz)

Controls the frequencies that will be passed to the effect.

Dry/wet mix

Controls the amount of distioning mixed into the output.

2.82 Rate shifter (rateShifter, 1417)

Stretches or compresses the input with a ringbuffer. Because of the ringbuffer you will get stretches of silence or clicks when the read pointer passes the write pointer. The ringbuffer is about 2.7-3.0s long, depending on the sample rate. Versions with variable buffer sizes or declipping code would be easy (but a bit less efficient); shout if you would find them useful.

Rate

The rate of the output signal; eg. 2.0 will double the speed. Negative numbers will play backwards. Pretty much any value will work, but the ranges give what most people are going to want to use. You can get some interesting sounds with very high numbers (e.g. 2000).

2.83 Retro Flanger (retroFlange, 1208)

A model of someone flanging the input. Models the tape saturation effects, and frequency smear of a manual flanger. The results are a slightly distorted, but more subtle flanger sound that you get from a normal digital flanger.

Average stall (ms)

The average time difference between the two tapes, per stall

Flange frequency (Hz)

The rate the tape is stalled at.

2.84 Reverse Delay (5s max) (revdelay, 1605)

A reverse delay not really modelled on any existing one. You'll want to set the Cross-fade Samples parameter to something reasonably small (but more than 20) for most applications, but you can try larger values if you start getting clicking.

Dry Level (dB)

Controls the level of the dry input signal in dB's.

Wet Level (dB)

Controls the level of the delayed signal in dB's.

2.85 Ringmod with two inputs (ringmod_2i1o, 1188)

This is a simple 2 input ring modulator. It is important that the modulator input is bounded to (-1, +1), otherwise you will get rubbish on the output.

Input

This is the audio input.

Modulator

This is the modulator input.

2.86 Ringmod with LFO (ringmod_1i1o1l, 1189)

This is a simple ring modulator and LFO.

2.87 Barry's Satan Maximiser (satanMaximiser, 1408)

Formerly Stupid Compressor. Thanks to Matt Yee-King for the name. Compresses signals with a stupidly short attack and decay, infinite ratio and hard knee. Not really as a compressor, but good harsh (non-musical) distortion.

Decay time (samples)

Controls the envelope decay time.

Knee point (dB)

Controls the knee roll-off point, ie. the point above which the compression kicks in. 0 will have no effect, -90 will remove virtually all dynamic range.

2.88 SC1 (sc1, 1425)

An high quality, reasonably low CPU cost RMS compressor designed for musical work. It has controls for the compression point, compression ratio and knee softness.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the compressor will start to kick in.

Ratio (1:n)

The gain reduction ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Makeup gain (dB)

Controls the gain of the makeup input signal in dB's.

2.89 SC2 (sc2, 1426)

A compressor with sidechain. Based on the code for SC1.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the compressor will start to kick in.

Ratio (1:n)

The gain reduction ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Makeup gain (dB)

Controls the gain of the makeup input signal in dB's.

2.90 SC3 (sc3, 1427)

A stereo compressor with sidechain input. Based on the code for SC1.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the compressor will start to kick in.

Ratio (1:n)

The gain reduction ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Makeup gain (dB)

Controls the gain of the makeup input signal in dB's.

Chain balance

Controls the chain signal used, 0 = Left+right in, 1 = Sidechain.

2.91 SC4 (sc4, 1882)

A stereo compressor with variable envelope follower for RMS / peak behaviour. Based on the code for SC1.

RMS/peak

The balance between the RMS and peak envelope followers. RMS is generally better for subtle, musical compression and peak is better for heavier, fast compression and percussion.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the compressor will start to kick in.

Ratio (1:n)

The gain reduction ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Makeup gain (dB)

Controls the gain of the makeup input signal in dB's.

Amplitude (dB)

The level of the input signal, in decibels.

Gain reduction (dB)

The degree of gain reduction applied to the input signal, in decibels.

2.92 SC4 mono (sc4m, 1916)

A mono compressor with variable envelope follower for RMS / peak behaviour. Based on the code for SC4.

RMS/peak

The balance between the RMS and peak envelope followers. RMS is generally better for subtle, musical compression and peak is better for heavier, fast compression and percussion.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the compressor will start to kick in.

Ratio (1:n)

The gain reduction ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Makeup gain (dB)

Controls the gain of the makeup input signal in dB's.

Amplitude (dB)

The level of the input signal, in decibels.

Gain reduction (dB)

The degree of gain reduction applied to the input signal, in decibels.

2.93 SE4 (se4, 1883)

A stereo expander with variable envelope follower for RMS / peak behaviour. Based on the code for SC1.

RMS/peak

The balance between the RMS and peak envelope followers.

Attack time (ms)

The attack time in milliseconds.

Release time (ms)

The release time in milliseconds.

Threshold level (dB)

The point at which the expander will start to kick in.

Ratio (1:n)

The gain expansion ratio used when the signal level exceeds the threshold.

Knee radius (dB)

The distance from the threshold where the knee curve starts.

Attenuation (dB)

Controls the gain of the output signal in dB's. Used to correct for excessive amplitude caused by the extra dynamic range.

Amplitude (dB)

The level of the input signal, in decibels.

Gain expansion (dB)

The degree of gain expansion applied to the input signal, in decibels.

2.94 Wave shaper (shaper, 1187)

This plugin reshapes the wave by an exponential function, inspiration was taken from the Nord module of the same name. If you are getting rubbish out then it's probably because the host isn't using the input/output range hints, which are very important for this plugin.

Waveshape

Positive values have an expanding effect, and negative values have a compressing effect.

2.95 Signal sifter (sifter, 1210)

Sorts and mixes blocks of the input signal to give a "bumpy ramp" effect. Certain types of input will produce silence on the output (mostly ones with only low frequency components). This is a very odd effect, and doesn't really have any music applications, but can produce some interesting noises.

2.96 Sine + cosine oscillator (sinCos, 1881)

This is a simple oscillator that outputs sinewaves with a 90 degree phase shift between them. The current implementation is very inefficient, but I will improve it later.

Base frequency (Hz)

The base frequency of the output waves.

Pitch offset

The pitch offset of the output waves. Final oscillator frequency is $base + 2^pitch$.

2.97 Single band parametric (singlePara, 1203)

A single band of a parametric filter.

Gain (dB)

The attenuation/gain of the eq.

Frequency (Hz)

The centre frequency (ie. point of most/least attenuation). Beware of high values for Frequency and Bandwidth, if the high pitch ($Frequency * 2^{Bandwidth}$) goes over half the sample rate you will get aliasing. Note: if your host offers you a frequency range between 0 and 0.4 then it's not rendering the input parameter correctly, the input frequency will actually be that number multiplied by the sample rate (e.g. 44.1kHz).

Bandwidth (octaves)

The pitch difference from the centre before the attenuation has reached half the gain.

2.98 Sinus wavewrapper (sinusWavewrapper, 1198)

Produces an unusual distortion effect, for a more amp like tone, see the valve saturation plugin (section 2.107).

2.99 Smooth Decimator (smoothDecimate, 1414)

Resample rate

The rate at which the output signal will be resampled

Smoothing

The amount of smoothing on the output signal.

2.100 Mono to Stereo splitter (split, 1406)

Takes a mono input signal, and outputs it to both left and right channel, thus "stereophizing" it.

2.101 Step Demuxer (stepMuxer, 1212)

Inputs up to 8 signals and switches between them on the output when then signal on the clock input goes high. This plugin is untested, and may not work.

2.102 Surround matrix encoder (surroundEncoder, 1401)

I haven't been able to test this plugin, so there may be bugs. I have successfully tested the algorithm, but the implementation is suspect.

What does it do?

It allows you to encode four channels of sound into a stereo compatible stream that will be decoded by a Dolby¹ Surround/Pro-Logic decoder into Left, Right, Center and Surround signals. This is not a proper implementation of Pro-logic, there is no Dolby B processing done on the surround channel, which would help, but would be using Dolby intellectual property.

Caveats

Obviously you can't wedge four channels into two without loss, so something has to give. You will probably notice significant cross-talk between the channels, but the decoder should do cross-talk correction, which will help a lot. A side effect of this is that it will make left-right panning unusual, the sources will dwell near the left and right speakers and zip across the centre channel. Because of this it is only really possible to master surround recordings through a pre-logic decoder. Do not attempt to use a conventional 5 point surround multichannel setup, it behaves very differently. In addition to this, output from this process is not entirely mono compatible, in mono output the L, C and R will be preserved as per a stereo recording (centre will be mixed equally), but the surround channel will be totally lost. Careful gain control on the output is required, as the level of the output will be greater than the L and R inputs, but different to the sum amplitude of the input signals. Widely panned reverb fed to the L and R channels will often leak into the S channel, if this is not desired (often it is)

¹"Dolby" is a trademark of Dolby Laboratories.

reduce the width of the stereo image. The encoding will survive some processes (eg. copying to CD, MD etc.), but may not survive conversion to MP3 or recoding to tape with azimuth errors.

Legalese

This is a four channel matrix encoder, it happens to be compatible with Dolby Surround Pro-Logic. This implementation does not convey a license nor imply a right under any patent, or any other industrial or intellectual property right of Dolby Laboratories.

L

Left channel input. Can be treated as per normal stereo recoding, except that the speaker should be at -22.5° , rather than the normal stereo -30° .

R

Right channel input. As per left channel.

C

Center channel input. Will be directly in front of the listener, stereo and mono compatible.

S

Surround channel. Should sound from the rear speakers, may also leak into the left and right. Has slight delay and bandwidth reduction (cut below 100 Hz, and above 7 KHz) for leakage and noise reduction and enhanced psychoacoustic effects. Not mono compatible.

2.103 State Variable Filter (svf, 1214)

An oversampled state variable filter with a few tweaks. Quite a nice State Variable Filter, tends to be unstable with high resonance and Q values, but good when kept under control.

Filter type (0=none, 1=LP, 2=HP, 3=BP, 4=BR, 5=AP)

Select between no filtering, low-pass, high-pass, band-pass, band-reject and all-pass.

Filter freq

Cutoff frequency, beware of high values with low sample rates.

Filter Q

The filter's Q, or cutoff slope.

Filter resonance

The filter's resonance, sort of separate from Q but very related (implemented with feedback). Do not use with the bandpass mode.

2.104 Tape Delay Simulation (tapeDelay, 1211)

Correctly models the tape motion and some of the smear effect, there is no simulation for the head saturation yet, as I don't have a good model of it. When I get one I will add it. The way the tape accelerates and decelerates gives a nicer delay effect for many purposes.

2.105 Transient mangler (transient, 1206)**2.106 Triple band parametric with shelves (triplePara, 1204)**

Actually five bands of eq, but the first and last are locked to shelving filters. For details see the single band parametric (section 2.97).

2.107 Valve saturation (valve, 1209)

A model of valve (tube) distortion, lacking some of the harmonics you would get in a real tube amp, but sounds good nonetheless. Taken from Ragnar Bendiksen's thesis: <http://www.notam02.no/~rbendiks/Diplom/Innhold.html>.

Distortion level

How hard the signal is driven against the limit of the amplifier.

Distortion character

The hardness of the sound, low for soft, high for hard.

2.108 Valve rectifier (valveRect, 1405)

Sag level

The level of power supply sag that will be caused by attacks.

Distortion

How harsh the distortion caused by the sag will be.

2.109 VyNil (Vinyl Effect) (vynil, 1905)

Year

The date of the recording/playback equipment to be simulated.

RPM

The rotational speed of the platter.

Surface warping

The degree of variation in height of the record surface.

Crackle

The number of scratches on the record surface.

Wear

The amount of wear on the grooves.

2.110 Wave Terrain Oscillator (waveTerrain, 1412)

A Wave Terrain oscillator, taken from Curtis Roads' example in *The Computer Music Tutorial*. Inputs x and y move the cursor around on a 2D landscape "wavetable" that is used to generate the output. The function used is $z = (x - y) * (x - 1) * (x + 1) * (y - 1) * (y + 1)$.

2.111 Crossfade (xfade, 1915)

Crossfade

Controls the degree to which the inputs are mixed into the output. A value of -1 means that the output is just the A input, and a value of 1.0 means that it is just the B input.

2.112 Crossfade (4 outs) (xfade4, 1917)

Crossfade

Controls the degree to which the inputs are mixed into the output. A value of -1 means that the output is just the A input, and a value of 1.0 means that it is just the B input.

2.113 z-1 (zm1, 1428)

A plugin that implements the z^{-1} function (a single sample delay).

3 Licensing

All this code is available under the GNU Public Licence, see the file "COPYING" included with the source for more details.