TSL SERIES · DELAY

LAP TSL-A95

DIGITAL DELAY · V4.2.7 · USER MANUAL

How it works - and how it differs from a plugin delay. TSL-A95 is a 12-bit digital delay modelled part by part: pre-emphasis, a limiter and an anti-alias filter in front of a 12-bit converter, a 32k-word memory, reconstruction filter and de-emphasis after it. The delay time is RANGE (2 … 512 ms in octaves) × FINE (×0.5 … ×2). FINE changes the clock of the converter, so turning it while sound is in the memory shifts the pitch like tape; with HOLD the memory loops and FINE transposes it by up to an octave. The delayed signal is band-limited to about 10 kHz and 12 bits; the direct signal is clean. The limiter sits in the delay path after the pre-emphasis, so bright material lights LIMIT first and the repeats come back compressed and darker - FEEDBACK returns them through the limiter and filter again, PHASE inverts the feedback. MODULATION modulates the clock (a vibrato of the delayed sound); its depth in cents grows with RANGE. BYPASS mutes the delayed path only; POWER switches the unit off. HI drives the input 6 dB harder.

A mastering plugin. TSL-A95 - like every plugin of the TSL series - solves the analogue board as a circuit in real time, and one instance takes a sizeable part of a CPU core (about 40 % of real time at 48 kHz on the development machine at the default 2× oversampling). It is made for the master bus, for stems and for a few key tracks, not as a send effect on every track of a session: a handful of instances is fine, a dozen will run the CPU out. Use it where it counts.

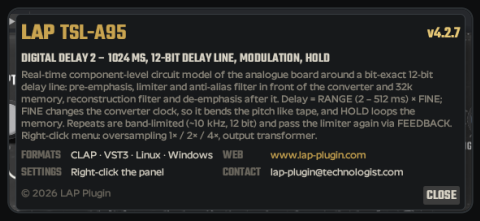


1 Working the panel

Action	Result
drag a knob	vertical drag changes the value; hold `Shift` for fine control. RANGE is a 9-position switch and moves one position at a time; the mouse wheel steps it one notch per click.
hover a knob	the exact value is printed below the knob scale while the pointer is over the knob
double-click a knob	resets to default (`Ctrl` -click does the same)

click HOLD	freezes the memory (amber LED lit): the stored sound repeats until HOLD is released, see Delay time, FINE and HOLD
click BYPASS	EFFECT OFF: mutes the delayed path (lit red bar = delayed path muted); the direct signal keeps going through the unit, so this is the unit's own effect switch, not a plugin bypass - that is POWER
sliders HI / LO, POWER, PHASE	click a slider to flip it: SENSITIVITY under INPUT, POWER under the name, PHASE under FEEDBACK
click the LAP logo	opens a small About window with the plugin version, formats and links
right-click the panel	the settings menu, see below
preset bar	above the panel, see Presets





Left: the settings menu (right-click the panel); the active choice of every row is highlighted. Right: the About window (click the LAP logo) with the version, formats and contact.

Menu row	Choices	What it does
SCALE	80 · 100	window size in percent: 1052 × 176, **1315 × 220** (default), plus the 30 px preset bar
STEREO	MIX · INV. MIX R	what the right output gets, see Stereo and levels; default **MIX**
OVERSAMPLING	1x · 2x · 4x	see Oversampling, latency and CPU; default **2x**
TRAFO	OFF · ON	the output transformer, see Stereo, levels and the output transformer; default **OFF**
SATURATION dB	−12 · 0 · +6 · +12 · +18	drive of the output transformer, see Stereo, levels and the output transformer; default **+6**
LAP	ABOUT	the About window

A click on SCALE, OVERSAMPLING or ABOUT closes the menu; STEREO, TRAFO and SATURATION keep it open, so the settings can be compared while playing.

The panel shares its frame with the other plugins of the TSL series. From the left: INPUT (dark grey cap, with the ACTIV and LIMIT LEDs above it and the SENSITIVITY HI / LO slider below), RANGE and FINE under the DELAY TIME bracket, DEPTH and SPEED with the HOLD button under the MODULATION bracket; the centre block with the LAP logo, the name and the POWER slider; on the right FEEDBACK

(PHASE slider below), MIX (blue cap) and OUTPUT (red cap) with BYPASS under the OUTPUT bracket. The whole panel is a 3D model rendered with ray-traced light, at three times the resolution of the default window so it stays sharp when enlarged or on high-DPI screens.

2 Presets



Element	What it does
category ▾	choosing a category loads its first preset at once. **Init** is the default state (INPUT max, LO, RANGE 4 ms, FINE ×1, DEPTH 0, SPEED about 2 Hz, HOLD off, FEEDBACK 0, PHASE off, MIX 50 %, OUTPUT max, POWER on, STEREO MIX, OVERSAMPLING 2x, TRAFO off, SATURATION +6 dB); **User** holds presets saved without a category; every other category is a folder of your own.
preset ▾	the presets of the chosen category; user presets are marked with a dot. "-" means the current settings do not come from a preset (a knob was moved).
◀ ▶	step through all presets - quick A/B between saved settings
Save	overwrites the selected user preset with the current settings
Save as...	name and category (type a new one or pick from ▾); OK saves, Cancel returns
Delete	deletes the selected user preset - click twice ("Really delete?")

A preset stores every knob and slider, HOLD, POWER and the settings menu (STEREO, OVERSAMPLING, TRAFO, SATURATION) - but not the window size (SCALE). It does **not** change BYPASS - that belongs to the session. Loading a preset moves the controls like a hand on the knobs: the audio keeps running, and the host records the change as automation if you are writing it.

Presets are plain files: Windows `%APPDATA%\LAP\TSL-A95\presets\`, Linux `~/.local/share/LAP/TSL-A95/presets/` - one JSON file per preset, one sub-folder per category. They can be copied between machines and systems.

3 Controls





Control	Range	Notes
INPUT	MIN … MAX	input level pot in front of the delay line; MAX is the nominal setting (−18 dBFS = nominal input level of the unit). Turn it down when LIMIT lights on everything.
SENSITIVITY	LO / HI	HI drives the input 6 dB harder into the limiter - for quiet sources or for a more compressed, darker repeat on purpose; the change is ramped over 150 ms. Default LO.
RANGE	2 · 4 · 8 · … · 512 ms	length of the memory, in octaves (9 positions); default 4 ms. With FINE it sets the delay time.
FINE	×0.5 … ×2	multiplies the delay time by changing the clock of the converter - continuous, and it bends the pitch of what is already in the memory while you turn it. Default ×1.
DEPTH	0 … 100 %	depth of the clock modulation; the law of the knob is steep (the middle of the travel is 10 %), because useful vibrato lives in the lower part. Default 0.
SPEED	10 … 0.1 Hz	rate of the modulation oscillator; the scale runs the way the unit prints it, fast on the left. Default about 2 Hz.
HOLD	button + LED	stops writing into the memory: its contents loop endlessly, FINE transposes the loop, DEPTH / SPEED wobble it
FEEDBACK	0 … 100 %	how much of the delayed signal goes back into the line - in front of the filter and the limiter, so every repeat is limited and filtered again. 100 % is just below self-oscillation. Default 0.
PHASE	ON / OFF	ON inverts the polarity of the feedback: at short times a different comb, at long times a repeat that alternates in polarity; the level of the repeats drops by about 6 dB. Default OFF.
MIX	ORIG … DELAY	one knob from the direct signal to the delayed one; the middle is an equal mix. Default 50 %.
OUTPUT	MIN … MAX	output level pot; MAX is the nominal setting. Default MAX.
BYPASS	ON / OFF	mutes the delayed path (EFFECT OFF); the direct signal still passes through the unit
POWER	ON / OFF	OFF switches the unit off: the dry signal passes, delay-compensated, and the plugin stops computing. ON again starts with an empty memory. Default ON.

4 Delay time, FINE and HOLD

The delay time is the length of the memory divided by the clock: **RANGE × FINE**. RANGE selects the memory length in octaves from 2 to 512 ms (at FINE ×1); FINE runs the converter clock from twice as fast (×0.5) to half as fast (×2), so the full span is 1 … 1024 ms. The clock sets the sample rate of the delay line as well - from 128 kHz at ×0.5 to 32 kHz at ×2 - and with it the resolution of the repeats, but the anti-alias and reconstruction filters are fixed, so the bandwidth of the repeat stays at about 10 kHz whatever the setting.

RANGE	2	4	8	16	32	64	128	256	512 ms
FINE ×0.5	1	2	4	8	16	32	64	128	256
FINE ×1	2	4	8	16	32	64	128	256	512
FINE ×2	4	8	16	32	64	128	256	512	1024

Delay time in milliseconds. The hover read-out of RANGE and FINE shows the resulting time.

FINE bends the pitch. Because FINE changes the clock, the sound already stored in the memory is read out faster or slower while you turn the knob - the repeats glide in pitch like a tape machine changing speed, and settle at the original pitch once the knob stops. Short RANGE settings with FEEDBACK (comb and flanger territory) are where this is most audible; long settings give the classic varispeed echo.

HOLD stops the write into the memory and loops what is in it - without any cross-fade, exactly as the unit does, so the seam of the loop is audible on sustained material and clean on rhythmic material. With HOLD on, FINE transposes the loop by up to an octave down (×2) or up (×0.5): a 1 kHz tone comes back at 500 Hz or at 2 kHz. DEPTH and SPEED modulate the loop like the live signal.

5 The sound of the delay line

The delayed signal passes the whole analogue chain of the unit and comes back different from the direct one - that is the point of this plugin. The direct signal (ORIG) is clean, 20 Hz – 20 kHz.

Pre-emphasis and limiter. In front of the converter the treble is raised (+14.5 dB at 10 kHz) and a limiter keeps the converter from clipping; after the memory the de-emphasis takes the treble down again. The limiter sees the pre-emphasised signal, so bright material drives it first - a hi-hat lights LIMIT long before a bass does. A limited repeat comes back compressed and darker; this grows with every pass through FEEDBACK. Use INPUT, or LO instead of HI, to stay below the limiter; use HI, or a hot INPUT, to lean into it.

Bandwidth and resolution. The anti-alias filter is fixed at −1 dB at 10 kHz and −38 dB at 20 kHz, so the repeats are band-limited to about 20 Hz – 10 kHz (+1 / −3 dB) whatever the delay time; the converter has 12 bits, about 60 dB of dynamic range. The quantisation is part of the model: the repeats carry the grain of a 12-bit line.

FEEDBACK returns the delayed signal in front of the anti-alias filter and the limiter, so each repeat is filtered and limited once more: a feedback tail loses treble and dynamics as it decays. At 100 % the loop sits just below self-oscillation, trimmed as the unit is on the bench. **PHASE** inverts the returned signal.

LEDs. Green **ACTIV** lights when a signal is present at the input of the delay line; red **LIMIT** lights while the limiter is working; amber **HOLD** shows the frozen memory. ACTIV and LIMIT fade over about 0.1 s so short peaks stay visible.

6 Modulation

DEPTH and SPEED modulate the clock of the converter, not a read pointer: the delay line speeds up and slows down like a tape transport with wow, and the delayed signal rises and falls in pitch. The waveform of the modulation oscillator is the one of the circuit - exponential rise and fall rather than a sine - and the excursion in cents grows with RANGE, because a longer memory is read for a longer time at the changed clock: at 4 ms a touch of chorus, at 512 ms a deep seasick wobble, for the same DEPTH. SPEED runs from 10 Hz (left) to 0.1 Hz (right). The DEPTH knob reaches 10 % at the middle of its travel; the upper half is for effects.

7 Stereo, levels and the output transformer



The unit is mono. A stereo track is summed to mono in front of the unit; MIX blends the direct and delayed signal and the result goes to both outputs. With **INV. MIX R** (settings menu, STEREO) the right output gets the delayed signal with inverted polarity - the second output of the unit - which puts the direct signal in the centre and spreads the repeats.

Levels: −18 dBFS corresponds to the nominal input of the unit with INPUT at MAX and LO, and to the nominal output with OUTPUT at MAX; the plugin passes the direct signal at unity. HI adds 6 dB in front of the limiter only.

Output transformer (TRAFO, default OFF). The unit itself has an electronic output; as in the other plugins of the TSL series, the signal can pass through a model of a real audio output transformer after it. The core with its magnetisation and hysteresis is solved sample by sample (at 4 × 44.1 / 48 kHz, i.e. 176.4 / 192 kHz whatever the session rate), driven from the output and loaded like a line input. It adds weight in the low end and low-order harmonics that grow with the bass content and the level, while mid and high frequencies stay clean; the level drops by about 1.5 dB, the same as in TSL-E26. Switching it on or off while playing is click-free and the host is told about the new latency.

SATURATION sets how hard the transformer is driven: the signal is raised by the chosen amount in front of the transformer and lowered by the same amount after it, so the level stays the same and only the core works harder or softer. Up to +12 dB the transformer adds only a trace of colour; at +18 dB the deep bass of loud passages goes into saturation. Changes of SATURATION are smoothed over 20 ms and can be made while playing.

8 Oversampling, latency and CPU

The analogue board is solved at a multiple of the session rate, counted from 44.1 / 48 kHz: **2×** (the default) runs it at 88.2 / 96 kHz in a 44.1 / 48 kHz session and at the session rate in 88.2 / 96 kHz sessions and above; 4× runs it at 176.4 / 192 kHz; 1× runs it at the session rate. The filters of the board are pre-warped and the feedback trimmer of the unit is set for the rate the board actually runs at, so FEEDBACK at MAX stays just short of self-oscillation at every setting. 4× makes the limiter and the clipping of the board a little smoother when you drive it hard; 1× saves CPU. The board never runs below 44.1 kHz (in a 32 kHz session 1× is 2×).

Session	1×	2×	4×	+ transformer
44.1 / 48 kHz	0	30	36	+36
88.2 / 96 kHz	0	0	30	+30
176.4 / 192 kHz	0	0	0	+0

Latency in samples of the session, reported to the host; it changes when OVERSAMPLING or TRAFO is switched. Switching either while playing is safe: the new unit is built in the background and cross-faded in (the repeats already in the line of the old unit fade out with it). The latency is the same with POWER off - the dry signal is delay-matched - so switching POWER while playing does not move anything.

Load. Stereo instance at 48 kHz on the development machine: about 20 % of real time at 1×, about 35 – 40 % at 2× (the default) and about 70 % at 4×; the transformer adds about 10 %. The load hardly depends on the settings (feedback and modulation change it by a few percent); with POWER off the plugin costs nothing. The analogue board - input stage, pre-emphasis, limiter with its detector, anti-alias and reconstruction filters, de-emphasis, mix and output stage - is solved as a circuit at every sample; the digital line (converter, memory, clock) is bit-exact and cheap.

9 Signal flow

Stage	What happens
input	stereo summed to mono · SENSITIVITY HI (+6 dB) / LO · INPUT level pot
pre-emphasis · limiter	+14.5 dB at 10 kHz · limiter with its detector (LIMIT LED), ACTIV detector
anti-alias filter	fixed low-pass, −1 dB at 10 kHz, −38 dB at 20 kHz; FEEDBACK (with PHASE) is summed in front of it
delay line	12-bit converter, 32k-word memory (length = RANGE, HOLD stops the write), converter; clock = FINE × modulation (DEPTH, SPEED), 32 … 128 kHz
reconstruction · de-emphasis	fixed low-pass and de-emphasis; BYPASS mutes this path
MIX · OUTPUT	ORIG … DELAY blend, OUTPUT level pot, both outputs (or inverted delay on the right)
output transformer	optional (TRAFO), per channel, with SATURATION; resampler back to the session rate