HUSH M4 mixer

Four-channel analogue club mixer, in development. Controls, signal flow, design targets, calibration and safety, as published at zakkhutoretsky.com/mixer.

IN DEVELOPMENT · REV 2027.4 (DRAFT) · DESIGN TARGETS, NOT MEASUREMENTS

> HUSH M4 is in development. This draft is published on zakkhutoretsky.com; it has no price and no ordering. Every figure marked **design target** is a goal for the finished product. Every figure marked **calculated** comes from the formulas shown, with nominal component values and ideal op-amps. **No value on this page was measured on a built unit**; measured results will be published separately, with their method. The drawings are original; no manufacturer schematic or branding was copied. Interviews are cited as the design brief, not as an endorsement.

DOCUMENT	REVISION	DATE
HUSH M4 Technical Manual	2027.4, draft	10 November 2026

STATUS

In development. Design targets, no measurements.

CONTENTS

1 Safety
2 About this manual
3 Controls and connections
4 Signal flow
5 Circuit description
6 Specification (design targets)
7 Calibration and alignment
8 Measurement plan
9 Care and handling
10 Revision history
11 References
Appendix A: derivations
Appendix B: schematics
Appendix C: design rationale and DVS1 sources
Appendix D: browser model audio specification (simulation)
Appendix E: browser app, DJ tools, sessions and offline (simulation)

1 SAFETY

This is a draft document for a product in development. Do not build mains-powered equipment from it without a qualified engineer's review and the safety approvals your country requires.

HAZARD	WHAT TO KNOW	DESIGN RESPONSE (DRAFT)
Mains voltage	Lethal. The primary side (inlet, fuse, switch, transformer) is the only part connected to mains.	Class I (earthed) enclosure; protective earth bonded to chassis and transformer core; fused live; double-insulated transformer; no user-accessible primary wiring. Design review against IEC 62368-1[R10] before any build.
Stored energy	The 4700 µF reservoirs hold charge after power-off.	Bleeder resistors across each reservoir (value to be set at prototype); wait and measure before opening.
Hot parts	Each regulator dissipates about 5.7 W.	Heatsinks inside the chassis, away from user access; vented enclosure.
Hearing	The headphone stage can deliver 2.83 Vrms into 32 Ω, far louder than comfortable.	Start with the PHONES knob at minimum. Do not wear headphones when connecting or powering up. Limit exposure.
Speaker damage	A DC fault or switch-on thump into a club system can damage drivers.	Series coupling capacitors on outputs, muting relay on the master outputs (prototype task), always power the mixer before the amplifiers.
Handling	A four-channel, transformer-powered mixer is heavy.	Lift with two people; a mass target is not set in this draft.

Safety standards are cited by title and catalogue page only; this draft does not claim compliance with any of them.

2 ABOUT THIS MANUAL

HUSH M4 is a four-channel analogue club mixer in development. The design brief started from what DVS1 has said in published interviews about how he plays (two CDJs and two Technics decks, an EQ that can take a band out, filters, no effects, booth monitoring, long sets). Appendix C maps each tailored feature to a cited quotation and lists where the design departs from him or nothing is documented. The interviews are cited as the design brief; this manual does not claim an endorsement.

Three words are used strictly. **Design target**: a goal the finished unit would have to meet. **Calculated**: worked out here from the cited formula with nominal component values and ideal op-amps. **Measured**: taken from a real unit with a stated method — **there are none in this document**. When a unit exists, a measurement report will follow the method of IEC 60268-3[R9].

The playable mixer on the web page is a software model built with the Web Audio API[R11]. Its filters and effects are not the analogue circuits in this manual, and nothing it plays is a measurement. Where the browser supports it, the model runs in an AudioWorklet with 32-bit float processing at the device's own sample rate, a four-band EQ and a split filter with exact bypass, optional texture, oversampled soft saturation that is transparent below 0 dBFS and a true-peak limiter on the output only, and it can play lossless tracks the user loads from their own device and record the master to a 32-bit float WAV file or a dithered 24-bit one. Appendix D lists what was measured on that code in a simulation. The hardware design has no limiter or compressor in the master or booth path.

3 CONTROLS AND CONNECTIONS

Channel strip (×4, identical)

CONTROL	RANGE / ACTION	NOTES
Input select PH / LN	Switches the strip between phono (RIAA, S1) and line (S2)	Four phono stages: any deck can go on any channel
TRIM	−15 dB to +15 dB on line (-14.9 to +14.3 calculated)	Set so the level LEDs peak in the amber
HI / HI MID / LO MID / LO	HI and LO: full kill to +6 dB. HI MID and LO MID: −27 dB to +10 dB. Centre (12:00) = exactly 0 dB, with a centre detent	HI shelf 3 kHz, HI MID bell 1.1 kHz, LO MID bell 350 Hz, LO shelf 180 Hz (S13). In the browser model: double-click or press Delete / Enter to reset
FILTER	Counter-clockwise: low-pass sweep. Clockwise: high-pass sweep. Centre dead zone: filter out of the circuit (relay true bypass)	12 dB/oct, low resonance (Q 0.71), no peak
CUE	Sends the strip pre-fader to the cue bus	Latching, lit when on
Channel fader	60 mm; sets the VCA control voltage	Replaceable. Audio does not pass through the fader.
Level LEDs	11 segments per channel (green, amber, red)	Read the post-fader signal

Master section

CONTROL	RANGE / ACTION	NOTES
MASTER ISOLATOR LOW / MID / HI	Full kill to +6 dB on the whole mix	Split and band amps as in S3b; the channel EQ (S13) reuses the split at 180 Hz
MASTER	Master output level	Feeds master XLR and the record RCA
BOOTH	Booth output level	Own balanced driver
CUE / MIX	Blend between cue bus and master in the headphones	SPLIT CUE puts cue in one ear and master in the other
PHONES	Headphone level	6.35 mm and 3.5 mm jacks, usable at the same time
TEXTURE / MODE	Amount (zero is off) and a mode button with LED: Clean SVF, Ladder, Comb, Bitcrush	Optional section, S11
FX SEND 1–4 / RETURN	Per-channel send to the spatial FX bus; one return level (zero is off); mode, time / size, feedback / decay and tone	Optional section, S12
METER DIM / DARK	Dims meters and legends to low amber	For dark rooms

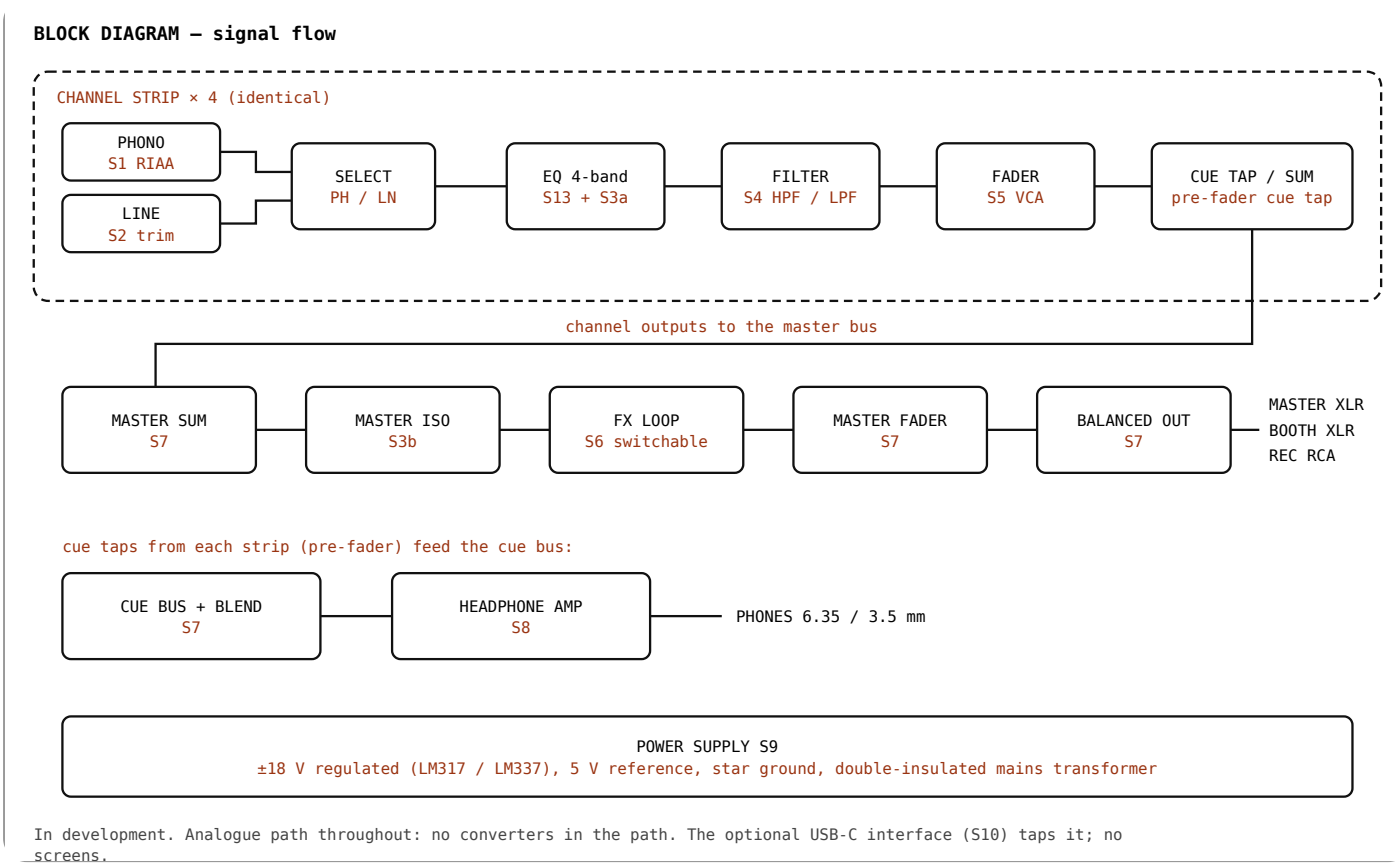
CONTROL	RANGE / ACTION	NOTES
MONO	Sums master to mono	For checking phase on mono systems
FX LOOP	Send level, return level, in / bypass	Out of circuit when off

Connections

CONNECTOR	TYPE	LEVEL / IMPEDANCE
Phono in ×4	RCA, unbalanced	MM, 47 kΩ, load capacitance switchable; chassis ground terminal
Line in ×4	RCA, unbalanced	Nominal −10 dBV to 0 dBV; trim ±15 dB
Master out	XLR, balanced	+28 dBu maximum (target), 95 Ω
Booth out	XLR, balanced	As master
Record out	RCA, unbalanced	Line level, fixed
FX send / return	6.35 mm, unbalanced	Nominal 0 dBV
Headphones	6.35 mm and 3.5 mm	250 mW into 32 Ω (target)
USB-C	USB Audio Class 2.0, optional	Record 4 channels plus master; playback to rear line outs (S10)
Power	Locking multi-pin input from the external linear supply	Mains inlet, fuse and protective earth are on the supply; chassis bonded to earth

4 SIGNAL FLOW

Each of the four strips runs input select, trim, four-band EQ, filter and fader, in that order, with a pre-fader cue tap. The strips sum into the master bus; the master isolator and the (switchable) FX loop follow, then the master fader and the balanced output drivers. The cue bus and headphone amplifier sit beside the master path. A single linear supply feeds everything. The whole path is analogue.



Block diagram

5 CIRCUIT DESCRIPTION

One short table per stage. The full schematics are in Appendix B and on /mixer/schematics.

S1 RIAA phono preamp (moving magnet)

ITEM	VALUE / FORMULA
Standard	Playback curve with time constants 3180, 318 and 75 µs[R1, R2]
U1	Non-inverting, G = 1 + Rf1/Rg1 = 1 + 2.21k/220 = 11 (+20.9 dB); Cg1 = 470 µF puts the stage high-pass corner at 1/(2π·220·470 µF) = 1.5 Hz
U2	Shelf: pole at Rf2·Cf2 = 20k × 160 nF = 3.2 ms (3180 µs), zero at Rf2·Cf2/(1 + Rf2/Rg2) = 318 µs. Rf2/Rg2 = 20k/2.21k ≈ 9
R3·C3	Passive low-pass, 5.76k × 13 nF = 74.9 µs (75 µs)
U3	Flat gain 1 + 6.19k/1k = 7.19 (+17.1 dB): total at 1 kHz = 38.1 dB, a design target of 38 dB for a 5 mV moving-magnet cartridge
Noise	U1 should be the lowest-noise op-amp in the design (NE5534-class, 3.5 nV/√Hz typ)[R7, R6]; the 220 Ω gain resistor contributes about 1.9 nV/√Hz (√(4kTR))
Result	Calculated RIAA error with these E96/E24 values and ideal op-amps: **±0.017 dB**, 20 Hz to 20 kHz. A real build adds op-amp bandwidth error, part tolerance and cartridge loading; the 20 Hz IEC rumble corner is not included

S2 Line input and trim

ITEM	VALUE / FORMULA
Trim	Inverting stage, G = −Rf/Rin, Rin = 10k, Rf = 1.8k (fixed) + P1 50k: -14.9 dB to +14.3 dB
Polarity	U5 is a unity inverter so the line path has the same absolute polarity as the phono path
Filters	RF low-pass 1/(2π·1k·220 pF) = 723 kHz; input coupling 2.2 µF into 47k = 1.5 Hz high-pass

S3a Isolator building block

ITEM	VALUE / FORMULA
Section	Second-order Butterworth Sallen-Key, equal resistors and equal capacitors, f0 = 1/(2πRC), K = 1 + Rf/Rg = 3 − 1/Q[R3]
LR4	Two cascaded sections give a Linkwitz-Riley 24 dB/oct filter, each stage Q = 0.71[R4]
Values	C = 10 nF. 300 Hz: R = 53.6k (f0 = 296.9 Hz). 3 kHz: R = 5.36k (f0 = 2969 Hz). Rg = 10k, Rf = 5.9k gives K = 1.59, Q = 0.709

S3b Three-band isolator (master section)

ITEM	VALUE / FORMULA
Split	LP4 at 300 Hz, HP4 at 300 Hz then LP4 at 3 kHz, HP4 at 3 kHz[R4]
Band gain	×2 band amp and a 10k linear pot: minimum = full kill, centre = unity (the pot divides by 2), maximum = +6 dB
Sum	Inverting sum, 10k per band. This is the master isolator. The channel EQ (S13) reuses the same split with the low crossover at 180 Hz and a fixed middle leg
Open task	Allpass delay trim between the three paths so the bands recombine flat in phase; to be done on a prototype

S4 Channel filter

ITEM	VALUE / FORMULA
Sections	Butterworth Sallen-Key (12 dB/oct), one low-pass and one high-pass, Q = 0.71[R3]
Low-pass	C = 4.7 nF, R = 1.87k + P5 500k dual-gang log: 18,108 Hz down to 67 Hz
High-pass	C = 10 nF, R = 2.67k + P6 250k dual-gang log: 63 Hz up to 5,961 Hz
Control	One knob: low-pass counter-clockwise, high-pass clockwise, with a centre dead zone
True bypass	A window comparator on a third wiper track drives relay K1 across the dead zone: IN goes straight to OUT and neither section is in the path. At the dead-zone edge both sections sit at their open ends (LP 18,108 Hz, HP 63 Hz)

ITEM	VALUE / FORMULA
Resonance	Butterworth, Q = 0.71: low resonance, no peak[R3]
Click	The relay switching click and the step at the dead-zone edge on the real circuit are prototype measurement items; nothing is claimed here

S5 Channel fader (control-voltage VCA)

ITEM	VALUE / FORMULA
Law	VCA control constant 6.1 mV/dB[R8]: 100 dB range needs 0.61 V
Control path	P7 10k linear fader on a 5 V reference → buffer → 47k / 6.49k divider → 0.607 V maximum
Audio path	Rin = Rf = 10k sets 0 dB at full travel; the audio never passes through the fader
Crossfader	Not fitted, as stated on the /mixer page (left out on purpose); a second control bus is a footprint option only

S6 FX send / return

ITEM	VALUE / FORMULA
Send	10k pot, follower, 10 µF and 100 Ω to the jack; nominal 0 dBV, unbalanced
Return	2.2 µF coupling, 47k bias, follower, 10k return pot, SW2 into the master mix
Bypass	SW2 removes the loop from the circuit entirely: no effect sits in the signal path when it is off

S7 Master, booth and cue

ITEM	VALUE / FORMULA
Sum	Inverting summer: gain per input = −Rf/Rin = −10k/10k = −1
Output	Impedance-balanced driver: follower plus unity inverter, 47.5 Ω in each leg, Rout = 2 × 47.5 = 95 Ω
Level	+28 dBu = 19.46 Vrms balanced, 13.76 V peak per leg. On ±18 V rails with 16.5 V usable swing per leg the balanced limit is about +29.6 dBu (1.6 dB above the +28 dBu design target). On ±15 V rails (13.5 V swing) it would be only +27.8 dBu: this is why the design uses ±18 V
Cue	Channel pre-fader taps into an inverting sum, then a cue / master blend pot into the headphone stage

S8 Headphone amplifier

ITEM	VALUE / FORMULA
Drive	P = V²/R. 250 mW into 32 Ω needs 2.83 Vrms (4.00 V pk, 88 mA rms); 250 mW into 64 Ω: 4.00 Vrms; 100 mW into 600 Ω: 7.75 Vrms (10.95 V pk)

ITEM	VALUE / FORMULA
Topology	Op-amp with a complementary emitter-follower stage inside the loop, 0.47 Ω emitter resistors, 10 Ω series output resistor[R12]
Damping	32 Ω / 10 Ω = 3.2: a design choice for short-circuit safety, not a high-damping design

S9 Power supply

ITEM	VALUE / FORMULA
Rails	±18 V, LM317 / LM337: Vout = 1.25 × (1 + R2/R1) = 1.25 × (1 + 3240/240) = 18.29 V[R5]
Rectifier	2 × 22 Vac, 50 VA toroid, bridge, 29.7 V DC nominal; 4700 µF reservoirs, ripple = I/(f·C) = 1.06 V pp at 0.5 A, 100 Hz
Headroom	10.4 V above the regulated output at the ripple trough, 7.3 V with mains −10% (LM317 asks for 3 V)[R5]
Heat	Each regulator dissipates about 5.7 W at 0.5 A: needs a heatsink
Budget	Estimated 124 op-amp sections (62 dual packages) × about 8 mA = 0.53 A per rail against the 0.5 A design current. An estimate, not a measurement

S10 USB-C interface (optional tap)

ITEM	VALUE / FORMULA
Purpose	Record the four channels and the master over one USB-C cable, and play host audio back to rear line outputs. A tap: it listens to the analogue path and is not part of it. The mixer works with the board removed or switched off
Converters	Two ES9842PRO 4-channel ADCs[R14]: eight record channels (CH1–4, master L and R, two spare). Datasheet figures: dynamic range +122 dB in 4-channel mode (+128 dB mono), THD+N -116 dB, PCM / TDM words up to 32 bit. The ES9843PRO is an alternative[R15]
USB device	XMOS XU316 running USB Audio Class 2.0, asynchronous, as in the XK-AUDIO-316-MC reference (8 in / 8 out, up to 384 kHz, 32-bit)[R16]. The playback DAC part is not chosen in this draft
32-bit float	A file and transport format. The converters output integer words; the host converts to 32-bit float when it writes the file. Float has a 24-bit significand (about 144 dB), which holds the 122 dB converter range (about 20.3 effective bits) without rounding. The real dynamic range is whatever the datasheet and the analogue noise floor give; float does not add range to the analogue path
Data rate	Eight channels of 32-bit words: 24.6 Mbit/s at 96 kHz, 49.2 Mbit/s at 192 kHz (calculated); USB 2.0 high speed carries 480 Mbit/s
Grounding	A digital isolator between the USB side and the audio ground (design target). Converters run from the board’s own audio clock
Not claimed	No measured figure for the assembled interface; the datasheet figures are the converter alone

S11 Texture (optional)

ITEM	VALUE / FORMULA
Placement	After S4, before the fader, switched by relay K2. Off is the analogue strip with no converter in it
Option A, analogue	Four-pole ladder or OTA filter with a drive stage in front, resonance by feedback, 24 dB per octave. The Moog ladder is analysed by Stilson and Smith[R17] and modelled non-linearly by Huovilainen[R18]. Component values are not set in this draft
Option B, DSP card	ADC, DSP and DAC per strip or shared and assignable. Modes: Clean SVF, Ladder (zero-delay-feedback form, after Zavalishin[R19]), Comb, Bitcrush; 4× oversampled (2× in the efficient mode); drive before the filter; resonance; LFO or envelope-follower modulation of the cutoff. Latency in the browser model: 51 samples
Controls	TEXTURE knob (amount; zero is off), MODE button with LED. The same signal flow runs in the browser model on /mixer
Not part of this revision	Formant and Noise / air modes

S12 Spatial FX send / return (optional)

ITEM	VALUE / FORMULA
Status	Design target. Nothing here is built or measured. The browser model on /mixer runs the same four modes in floating point so the idea can be heard; it is not the analogue design
Topology	Four post-fader taps through 10k resistors into an inverting sum (gain −Rf/Rin = −1), a send level pot, the FX engine, a return level pot and a hard-bypass switch SW3 into the master sum of S7. With SW3 out the engine is out of circuit, as for S6
Option A, BBD ensemble	MN3207, 1024 stages[R22]: delay td = 1024 / (2 × fclock); the datasheet gives 2.56 ms at 200 kHz to 51.2 ms at 10 kHz and a +4 to +10 V supply. A 6.4 ms voice needs 1024 / (2 × 0.0064) = 80 kHz (calculated). Typical THD 0.4 % and S/N 73 dB under the datasheet's own test conditions[R22]. Three voices with clocks 120 degrees apart give the ensemble; a clock generator and anti-alias / reconstruction filters are needed and are not drawn. Companding is a design choice for noise, with no part chosen
Option B, digital delay	PT2399[R21]: 44 kbit RAM, delay set by the VCO frequency through one external resistor (10 kΩ recommended, up to 50 kΩ allowed in the datasheet note), THD 0.5 % and SNR better than 90 dB typical, +5 V. Echo feedback would be an op-amp mix of the output back to the input with a pot; wow and flutter from modulating the VCO resistor current is a circuit idea of the designer, not a datasheet claim, and is untested. The datasheet's delay-time table is on a page not included in the copy read, so no delay range is claimed
Option C, spring and plate	Real springs and plates are not modelled by a short BBD. A DSP card (ADC, DSP, DAC) or a dedicated reverb IC would be needed; none is chosen. The plate algorithm in the browser model follows Dattorro[R20]: input diffusers 142, 107, 379, 277 samples, tanks 672 + excursion / 4453 / 1800 / 3720 and 908 + excursion / 4217 / 2656 / 3163 samples at 29,761 Hz, decay-diffusion 0.70 and 0.50, 0.6 output taps from the paper's Table 2. The spring is the designer's own: 30 first-order all-pass stages per line for dispersion plus a delay and damping loop; it is not taken from a source

ITEM	VALUE / FORMULA
Controls (model)	Per-channel send; time / size, feedback / decay, tone, return; mode; tap tempo or tempo sync for the echo. Feedback is capped inside the model: echo 0.85, spring 0.95, plate decay 0.93, BBD regeneration 0.5. Return at 0 adds exact zeros
Browser model limits	No hiss, clock noise or aliasing of a real BBD is modelled; tape echo has no head or tape model beyond a saturating soft clip and low-pass filters. CPU: measured in Node on the test runner only (see tests/dsp_test.js); no phone was measured

S13 Channel EQ (four-band, per channel)

ITEM	VALUE / FORMULA
Status	Design target. Nothing here is built or measured. The browser model on /mixer is floating-point DSP that follows the same control law; it is not a simulation of this circuit
Bands	HI shelf 3 kHz: +6 dB to full kill. HI MID bell 1.1 kHz: +10 dB to −27 dB. LO MID bell 350 Hz: +10 dB to −27 dB. LO shelf 180 Hz: +6 dB to full kill. Every knob is exactly 0 dB at 12:00 (centre detent). Where these figures come from is stated in the manual, section 5, S13
Strip order	Trim, LO MID bell, HI MID bell, three-way LR4 split with LO and HI band pots, then the filter (S4) and the fader (S5). The cue tap stays pre-fader
Bells (two stages)	Each bell is an inverting op-amp stage with a 10k linear pot and a band-pass filter tuned to the bell frequency, the virtual-earth boost / cut arrangement described by ESP[R23]. ESP lists multiple-feedback, twin-T and state-variable band-passes as options; this design picks multiple-feedback for part count (a design choice). At the pot centre the band-pass has no effect and the stage is a unity inverter, so it is flat. Two inverting stages in series keep absolute polarity. Bell Q about 0.9 at full boost (design choice). Component values are not calculated in this draft
Shelves (kill)	A plain Baxandall shelf cannot reach full kill, so LO and HI reuse the S3a LR4 split instead[R3, R4]: crossovers at 180 Hz (R = 88.7k, f0 = 179.4 Hz) and 3 kHz (R = 5.36k, f0 = 2969 Hz) with C = 10 nF. LO and HI bands each get a ×2 band amp and a 10k linear pot (minimum = kill, centre = unity, maximum = +6 dB); the middle leg gets ×2 then a fixed divide-by-2; an inverting sum recombines
Kill depth	Design target: beyond 60 dB in the stop band of the LR4 split, as for S3b. Not measured
Model versus circuit	The browser model uses second-order shelf and peaking sections (the textbook biquad forms, not cited to a page in this draft) with the same corner and centre frequencies and the same plateau levels, and a −80 dB floor for kill. The LR4 split has steeper transitions near 180 Hz and 3 kHz, so the two differ in the transition regions
Headroom	Boosts stack: with every band up the strip can add well over 10 dB at some frequencies. Headroom for that has not been analysed in this draft. The model on /mixer keeps its limiter on the master output only
Op-amp sections	Bells: 2 sections each, 4 per channel. Split: 12 sections per channel (see the budget in S9)

Where the S13 figures come from. The four corner and centre frequencies and the boost and cut ranges are the ones Allen & Heath publish for the stereo channels 1–4 of its Xone:96 mixer[R24] (LO 180 Hz, +6 dB to kill; LO MID 350 Hz, +10 / −27 dB; HI MID 1.1 kHz, +10 / −27 dB; HI 3 kHz, +6 dB to kill). They are used as a published reference for a familiar DJ EQ layout. The circuit, the shelf and bell shapes (shelf Q 0.71, bell Q 0.9), the knob design and the panel drawing are this design's own; the guide gives no EQ schematic, and no part of that maker's design, artwork or logo

was copied. No connection with, or endorsement by, Allen & Heath is implied. **Relation to DVS1:** he is quoted as saying he does not need a 4-band EQ (q6), so this four-band EQ departs from his statement; it was added at the site owner's request. The only related account is a writer reporting him ducking a high-mid on an Xone:96 (q26), a band a three-band isolator does not have. See Appendix C.

6 SPECIFICATION (DESIGN TARGETS)

These are goals, not measurements. The third column says whether a figure is a design target or a calculation from the formulas in section 5.

ITEM	VALUE	BASIS
Channels	4, each phono / line switchable	design target
Trim range (line)	-14.9 dB to +14.3 dB	calculated (G = −Rf/Rin)
Phono gain at 1 kHz	38.1 dB (5 mV MM → about 0.4 V)	calculated
Phono RIAA accuracy	±0.2 dB (20 Hz to 20 kHz)	design target
Phono RIAA error, nominal values, ideal op-amps	±0.017 dB	calculated, not a measurement
Phono input	47 kΩ with switchable loading (capacitance values set at prototype)	design target
Channel EQ	4-band (S13): HI shelf 3 kHz +6 dB to kill; HI MID bell 1.1 kHz +10 / −27 dB; LO MID bell 350 Hz +10 / −27 dB; LO shelf 180 Hz +6 dB to kill; exactly flat at 12:00. Shelf split R = 88.7k (179.4 Hz) and 5.36k (2969 Hz)	design target; split frequencies calculated
Master isolator	3-band, LR4 crossovers at 300 Hz and 3 kHz (actual 296.9 Hz and 2969 Hz); full kill to +6 dB	calculated / design target
Channel filter	Bipolar HPF / LPF, 12 dB/oct; LPF 18,108 Hz to 67 Hz, HPF 63 Hz to 5,961 Hz; true bypass at centre	calculated
Channel fader	60 mm conductive plastic, control-voltage VCA, 100 dB range; no crossfader	design target
Master	3-band master isolator, master and booth level, mono switch, VU meters L / R	design target
Maximum output (balanced)	+28 dBu; rail-limited about +29.6 dBu on ±18 V (1.6 dB margin)	design target / calculated
Output impedance (balanced)	95 Ω	calculated
Frequency response	10 Hz to 40 kHz, ±0.25 dB (line in to master out)	design target

ITEM	VALUE	BASIS
THD+N	below 0.003 % at +4 dBu, 1 kHz	design target
Signal-to-noise	above 105 dB A-weighted, line in to master out	design target
Crosstalk	below −95 dB at 1 kHz between channels	design target
Headphones	6.35 mm and 3.5 mm; 250 mW into 32 Ω (2.83 Vrms)	design target / calculated
Power	External linear supply, ±18 V; mains transformer 2 × 22 Vac, 50 VA	design target
Effects loop	Switchable send / return, hard-bypassed when off	design target
Audio path	Fully analogue, balanced out; no converters in the signal path	design target
Outputs	Master XLR, booth XLR, record RCA; transformer-isolated on request	design target
USB-C interface (optional)	Record 4 channels plus master and play back; USB Audio Class 2.0; converters on a tap off the analogue path (S10)	design target
Texture (optional)	Clean SVF, ladder, comb, bitcrush; resonance, drive, LFO / envelope; analogue ladder / OTA option and DSP card (S11)	design target
Spatial FX (optional)	Send / return with a BBD ensemble, digital delay or DSP card option; four browser-model modes: tape echo, spring, plate, BBD ensemble (S12)	design target
Left out on purpose	DJ software, screens, menus, crossfader	design target

Level and rail consistency

The +28 dBu maximum output equals 19.46 Vrms across the balanced pair, or 13.76 V peak per leg. With ±18 V rails and an assumed 16.5 V usable swing per leg, the driver can reach about +29.6 dBu, so 1.6 dB of margin. On ±15 V rails the same driver would top out near +27.8 dBu, which is why the supply is ±18 V and why the swing figure must be confirmed against the chosen op-amp's datasheet.

Consistency with the /mixer page

- The /mixer specification and this manual both describe op-amp stages: low-noise NE5534 / NE5532-class op-amps for the RIAA stages and an inverting op-amp summer. No discrete or Class A stage is claimed; a discrete input stage is an open item for a later revision.
- Both state no crossfader. The VCA fader stage could host one; it is not fitted in this design.
- The +28 dBu maximum output needs ±18 V rails, as shown above, not the more usual ±15 V.

7 CALIBRATION AND ALIGNMENT

A proposed procedure for a prototype. Nothing here has been carried out. Use an isolation transformer on the bench.

STEP	WHAT TO DO	TARGET
1 Supply	Power up through a series lamp or variac. Measure ± rails unloaded and loaded; check ripple with a scope.	+18.3 / −18.3 V ± 0.4 V; ripple below the section 5 figure
2 Offsets	With inputs shorted, measure DC at every stage output.	Within a few millivolts; trim or servo as needed
3 VCA	Set the fader to minimum and maximum, adjust the VCA offset and symmetry trims per the chosen VCA's datasheet, then check 0 dB at full travel.	0 dB at full travel (±0.2 dB)
4 RIAA	Inject through an inverse-RIAA network (or an instrument that produces one) and measure the response at 20, 50, 100, 500 Hz, 1, 2.1, 5, 10, 20 kHz, referenced to 1 kHz.	Within ±0.2 dB, tighter if the parts allow
5 Trim and line	Check the trim range at both ends and the polarity through S2.	About −15 dB to +15 dB; polarity unchanged vs. phono
6 Master isolator	Sweep the generator; confirm −6 dB at 300 Hz and 3 kHz on each band, and full kill at the pot minimum.	Crossovers within 5 % of design; kill beyond 60 dB
6b Channel EQ	On each strip, sweep the generator with every knob centred, then one knob at a time to each end. Check the bell centres, the shelf corners at 180 Hz and 3 kHz, the ranges (+6 / kill, +10 / −27 dB), and the flat response with all knobs at 12:00.	Centres and corners within 5 % of design; ranges within 1 dB; kill beyond 60 dB; flat within 0.5 dB at 12:00 (targets for a prototype)
7 Filter	Check the dead-zone relay bypass and the sweep ends of P5 and P6.	End frequencies within 15 % of the calculated range
8 Output stage	Set the headphone output stage idle current at the prototype (Vbe-multiplier trim); check for thermal runaway.	Idle current chosen at the bench; stable at 50 °C case
9 Balanced out	Check level balance between pins 2 and 3 and the common-mode rejection into a balanced load.	Legs within 0.1 dB of equal
10 Soak	Run for four hours at the maximum level into a dummy load.	No drift beyond the offsets tolerance; heatsinks within limits

8 MEASUREMENT PLAN

When a unit exists, the figures in section 6 would be measured following the methods of IEC 60268-3[R9] and reported with the test conditions (source impedance, load, bandwidth, weighting, reference level). Until then the table in section 6 remains a list of design targets. AES17, the method for digital audio equipment, would apply to the optional USB-C interface (S10). **Publication plan (method only):** every result would be published with its test conditions, the instrument and its settings, the number of units tested and the raw files, and a figure would appear on the product pages only after it has been measured on a real unit. No result is given here.

QUANTITY	CONDITION (PROPOSED)	INSTRUMENT
Frequency response	Line in to master out, 1 kHz reference, 10 Hz to 40 kHz	Audio analyser or swept generator + scope

QUANTITY	CONDITION (PROPOSED)	INSTRUMENT
THD+N	+4 dBu, 1 kHz, 22 Hz–22 kHz bandwidth and 80 kHz bandwidth	Audio analyser
Signal-to-noise	Ref +4 dBu, A-weighted, inputs terminated with 150 Ω (line) / 47 kΩ + 200 pF (phono)	Audio analyser
Crosstalk	Drive one channel, measure the others, 1 kHz and 10 kHz	Audio analyser
Maximum output	Clipping point into 600 Ω and 10 kΩ, balanced	Scope + analyser
Filter bypass	Switch the dead-zone relay with a steady tone, record the largest sample step against the tone's own step	Scope or recorder
USB-C record path	Loop a known level from a generator through the tap; compare file level and noise with the analogue measurement	Audio analyser + recording software
Texture	Each mode with a swept tone at three amounts; record level, noise and any aliasing products	Audio analyser

9 CARE AND HANDLING

Keep liquids away; the faders and pots are user-replaceable but the panel is not sealed. Clean with a dry cloth. Power the mixer on before the amplifiers and off after them. Do not block the ventilation slots. Replace the mains fuse only with the stated type.

10 REVISION HISTORY

REVISION	DATE	CHANGE
2027.4 draft	10 November 2026	The browser model becomes an installable web app: offline app shell, saved sessions, per-deck BPM, tap tempo, pitch fader, loops, hot cues, nudge and an optional sync, a track library, Web MIDI learn, keyboard shortcuts and a first-run tour. These are tools of the software model; the hardware design is unchanged and still has no tempo, BPM or sync function. Adds Appendix E. Departs from a DVS1 statement (q24); see Appendix C.
2027.3 draft	10 November 2026	Replaces the per-channel three-band isolator with a four-band channel EQ (S13, design target): HI shelf, HI MID bell, LO MID bell, LO shelf, exactly flat at 12:00. Adds the matching browser model, centre detent and reset on /mixer, and references R23 and R24. Adds lossless track decks to the browser model (WAV, AIFF, FLAC from the user's device, never uploaded). Reworks its audio path (native sample rate, 4× oversampling of the nonlinear stages, exact bypass, true-peak limiter at −1 dBTP, dithered 24-bit export) and adds Appendix D with the simulation results. The master isolator (S3b) is unchanged. Departs from a DVS1 statement (q6); see Appendix C.
2027.2 draft	10 November 2026	Adds the optional spatial FX send / return (S12, design target) and the matching browser model on /mixer. Texture section S11 and USB interface S10 unchanged.

REVISION	DATE	CHANGE
2027.1 draft	10 November 2026	First draft: controls, signal flow, design targets, circuits S1–S9, calibration plan, safety. Values from tools/mixer_values.py.

11 REFERENCES

[R1] RIAA equalization: time constants 3180, 318 and 75 µs (IEC 60098 / RIAA) — Wikipedia, RIAA equalization (https://en.wikipedia.org/wiki/RIAA_equalization)

[R2] S. P. Lipshitz, "On RIAA Equalization Networks", Journal of the Audio Engineering Society 27(6), pp. 458–481, June 1979 — AES E-Library record (https://aes.org/publications/elibrary-page/?id=3197)

[R3] J. Karki, "Analysis of the Sallen-Key Architecture", Texas Instruments application report SLOA024B — TI SLOA024B (https://www.ti.com/lit/an/sloa024b/sloa024b.pdf)

[R4] S. Linkwitz, "Active Filters": LR2 and LR4 crossovers as cascaded Sallen-Key sections, Q0 = 0.71 per LR4 stage — Linkwitz Lab, Active Filters (https://linkwitzlab.com/filters.htm)

[R5] Texas Instruments, LM317 3-Pin Adjustable Regulator datasheet (Rev. Z): $V_o = 1.25\ V \times (1 + R2/R1) + I_{adj} \cdot R2$ — TI LM317 datasheet (https://www.ti.com/lit/ds/symlink/lm317.pdf)

[R6] Texas Instruments, NE5532 dual low-noise operational amplifier datasheet (5 nV/√Hz typical at 1 kHz) — TI NE5532 datasheet (https://www.ti.com/lit/ds/symlink/ne5532.pdf)

[R7] Texas Instruments, NE5534 low-noise operational amplifier datasheet (3.5 nV/√Hz typical at 1 kHz) — TI NE5534 datasheet (https://www.ti.com/lit/ds/symlink/ne5534.pdf)

[R8] THAT Corporation, 2180 series voltage-controlled amplifier datasheet (control constant 6.1 mV/dB) — THAT Corporation datasheets (https://thatcorp.com/datasheet-dynamics-processors/)

[R9] IEC 60268-3:2018, Sound system equipment — Part 3: Amplifiers (specification and measurement methods) — IEC webstore (https://webstore.iec.ch/en/publication/32788)

[R10] IEC 62368-1:2023, Audio/video, information and communication technology equipment — Part 1: Safety requirements — IEC webstore (https://webstore.iec.ch/en/publication/69308)

[R11] W3C, Web Audio API (the browser demonstration on /mixer uses it; it is not the analogue design) — W3C Web Audio API (https://www.w3.org/TR/webaudio/)

[R14] ESS Technology, ES9842PRO 4-channel ADC datasheet v0.5.1 (dynamic range, THD+N, 32-bit PCM / TDM words) — ES9842PRO datasheet (https://www.esstech.com/wp-content/uploads/2025/10/ES9842PRO_DS_v0.5.1.pdf)

[R15] ESS Technology, ES9843PRO 4-channel ADC datasheet v0.4.1 (alternative part) — ES9843PRO datasheet (https://www.esstech.com/wp-content/uploads/2025/11/ES9843PRO_DS_v0.4.1.pdf)

[R16] XMOS, XK-AUDIO-316-MC USB audio multichannel reference: USB Audio Class 2.0, asynchronous, XU316 — XMOS sw_usb_audio, app_316_mc (https://github.com/xmos/sw_usb_audio/blob/develop/doc/rst/app_316_mc.rst)

[R17] T. Stilson and J. Smith, "Analyzing the Moog VCF with Considerations for Digital Implementation", ICMC 1996 — CCRMA, Stilson and Smith (https://ccrma.stanford.edu/~stilti/papers/moogvcf.pdf)

[R18] A. Huovilainen, "Non-linear digital implementation of the Moog ladder filter", DAFx-04 — DAFx-04 paper archive (https://www.dafx.de/paper-archive/2004/P_061.PDF)

[R20] J. Dattorro, "Effect Design, Part 1: Reverberator and Other Filters", Journal of the Audio Engineering Society 45(9), pp. 660–684, September 1997 (plate topology, delay lengths and output taps, Fig. 1 and Table 2) — CCRMA, Dattorro (https://ccrma.stanford.edu/~dattorro/EffectDesignPart1.pdf)

[R21] Princeton Technology Corp., PT2399 Echo Processor IC datasheet V1.9, March 2022 (44 kbit RAM, VCO delay set by an external resistor, THD 0.5 % typ, SNR better than 90 dB typ, +5 V application circuits) — PT2399 datasheet (https://www.princeton.com.tw/Portals/0/activeforums_Attach/PT2399-s.pdf)

[R22] Panasonic MN3207, 1024-stage low-voltage low-noise BBD datasheet (delay 2.56 to 51.2 ms for a 200 to 10 kHz clock, supply +4 to +10 V, THD 0.4 % typ, S/N 73 dB typ). Read from a copy hosted by Synthrotek. — MN3207 datasheet (hosted copy) (https://www.synthrotek.com/datasheets/MN3207.pdf)

[R23] R. Elliott (ESP), "Equalisers, The Various Types And How They Work", Elliott Sound Products, March 2015, updated February 2021 (Baxandall and James tone controls; boost / cut stages with a band-pass filter between virtual-earth op-amp inputs; multiple-feedback, twin-T and state-variable band-pass options; gyrators) — ESP, Equalisers (https://sound-au.com/articles/eq.htm)

[R24] Allen & Heath, Xone:96 User Guide, channel equaliser and hardware specification (stereo channels 1-4: HI 3 kHz +6 dB / kill, HI MID 1.1 kHz +10 / -27 dB, LO MID 350 Hz +10 / -27 dB, LO 180 Hz +6 dB / kill). Used only as a published reference for the control layout and ranges; no circuit is published there. — Allen & Heath support site (https://support.allen-heath.com/hc/en-gb/articles/43053562464017-Xone-96-User-Guide)

[R19] V. Zavalishin, The Art of VA Filter Design, Native Instruments. Cited by title and author; no link was checked for this draft.

[R12] D. Self, Small Signal Audio Design, 3rd edition, Focal Press / Routledge (2020). Book, no link.

[R13] P. Horowitz and W. Hill, The Art of Electronics, 3rd edition, Cambridge University Press (2015). Book, no link.

Interview and rider sources for the design brief are in Appendix C.

A DERIVATIONS AND COMPONENT VALUES

All values come from `tools/mixer_values.py`; running it prints the same numbers.

RIAA

Playback curve H(s) = (1 + s·318 µs) / ((1 + s·3180 µs)(1 + s·75 µs))[R1, R2]. U2 provides the 3180 / 318 µs shelf: for G = 1 + Zf/Rg with Zf = Rf in parallel with Cf, the pole is at Rf·Cf and the zero at Rf·Cf/(1 + Rf/Rg), so Rf·Cf = 3180 µs and Rf/Rg = 9 give a 318 µs zero. A passive R·C = 75 µs follows. A non-inverting stage cannot fall below unity gain, which is why the 75 µs roll-off is passive and the gain is made up in U3.

CHECK FREQUENCY	ERROR VS IDEAL RIAA (DB)
20	0.01
50	0.011
100	-0.001
500	-0.005
1000	-0.0
2122	0.006
5000	0.011
10000	0.013
20000	0.013

Sallen-Key and Linkwitz-Riley

For equal R and C: f0 = 1/(2πRC), Q = 1/(3 − K), K = 1 + Rf/Rg[R3]. K = 1.586 gives Q = 0.707 (Butterworth); two such sections in cascade make an LR4 filter with the −6 dB point at f0[R4]. With C = 10 nF: R(300 Hz) = 53.05 kΩ → 53.6k; R(3 kHz) = 5.305 kΩ → 5.36k.

Power

Vout = 1.25 × (1 + R2/R1) with R1 = 240 Ω and R2 = 3.24k gives 18.29 V[R5]. Reservoir ripple ΔV = I/(f·C) with f = 100 Hz for full-wave rectification.

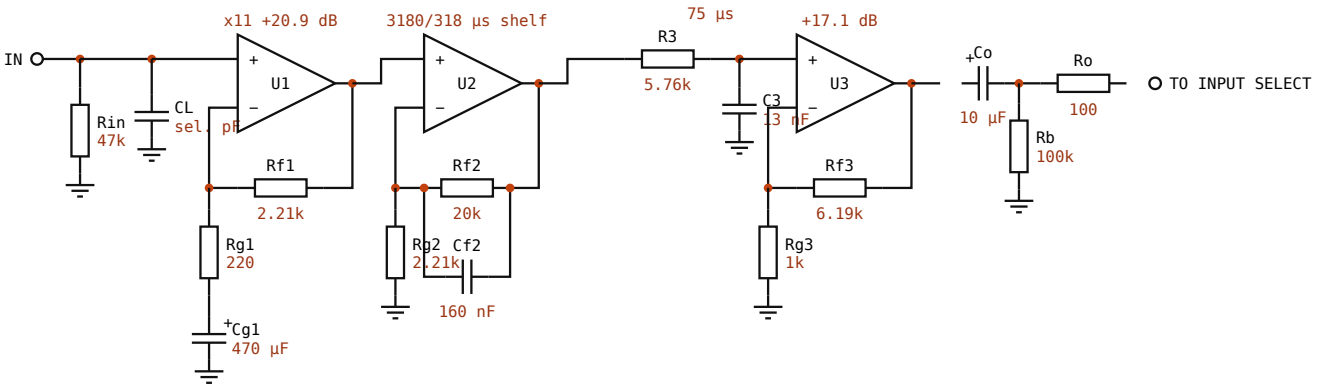
Levels

0 dBu = 0.7746 Vrms. +28 dBu = 0.7746 × 10^(28/20) = 19.46 Vrms. Balanced peak per leg = Vrms × √2 / 2 = 13.76 V.

B SCHEMATICS

Original drawings. Each is followed by its value table in section 5.

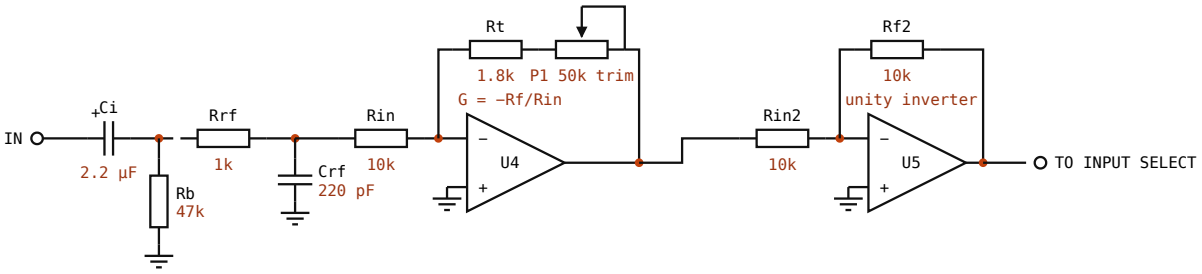
S1 RIAA PHONO PREAMP (MM) — one of four identical channels



Ideal-op-amp calc: RIAA error max 0.017 dB, 20 Hz to 20 kHz (E96 / E24 nominal values). 1 kHz gain 38.1 dB. U1: low-noise op-amp; U2, U3: dual low-noise op-amp. All op-amps on ±18 V.

S1 RIAA phono preamp (moving magnet)

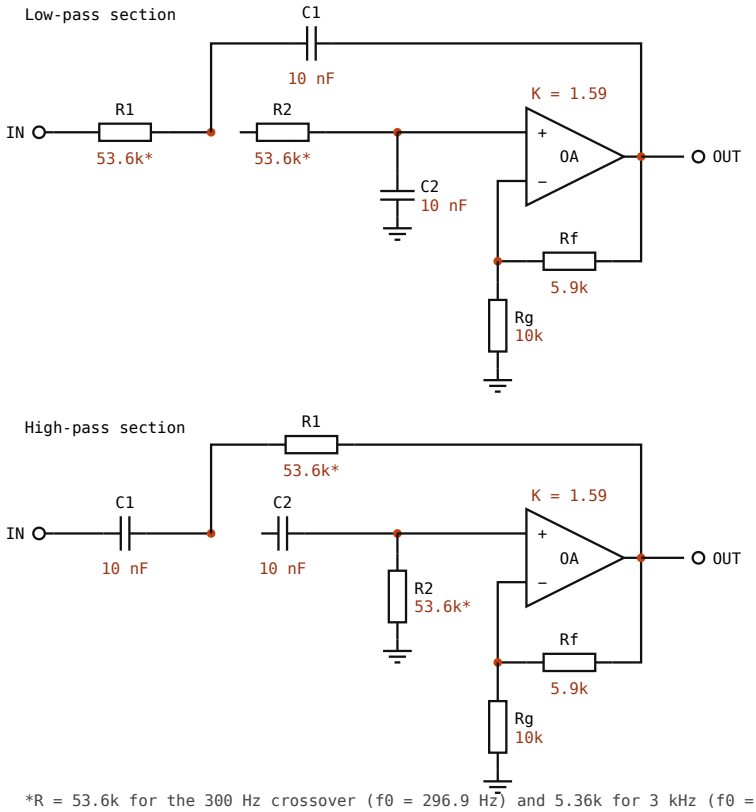
S2 LINE INPUT AND TRIM — one of four identical channels



Trim range -14.9 dB to +14.3 dB (formula G = −Rf/Rin, Rf = 1.8k to 51.8k, Rin = 10k). RF filter corner 723,432 Hz (1/2πRC), input high-pass 1.5 Hz.

S2 Line input and trim

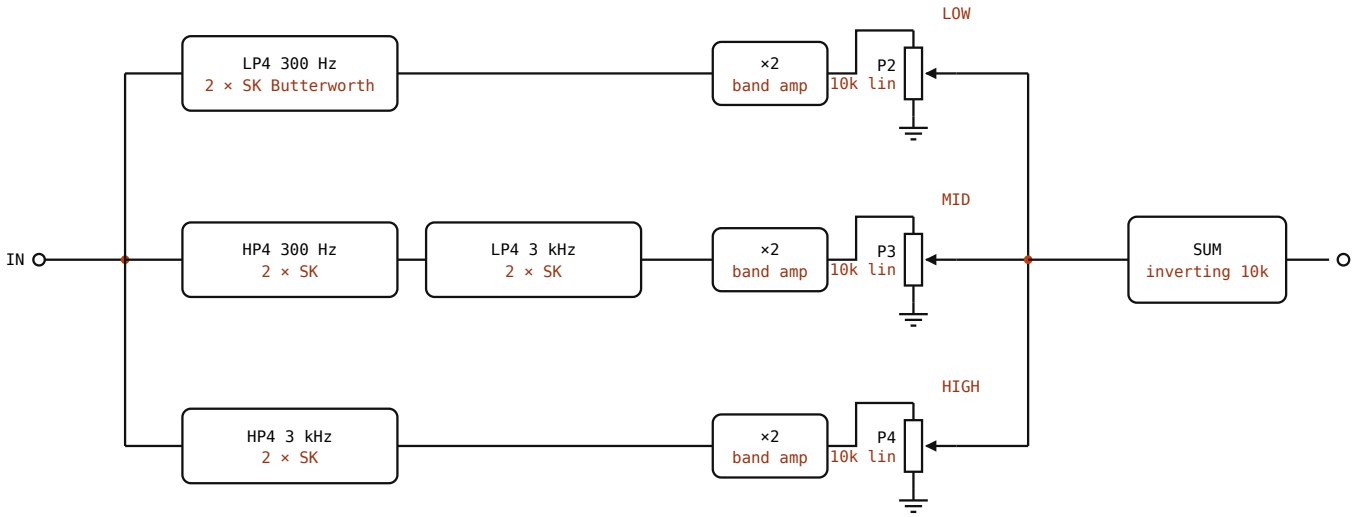
S3a ISOLATOR BUILDING BLOCK — 2nd-order Butterworth Sallen-Key (two in cascade = LR4)



*R = 53.6k for the 300 Hz crossover (f0 = 296.9 Hz) and 5.36k for 3 kHz (f0 = 2969 Hz); C = 10 nF; K = 1 + Rf/Rg = 1.590, Q = 1/(3 − K) = 0.709 (Butterworth: 0.707).

S3a Isolator building block

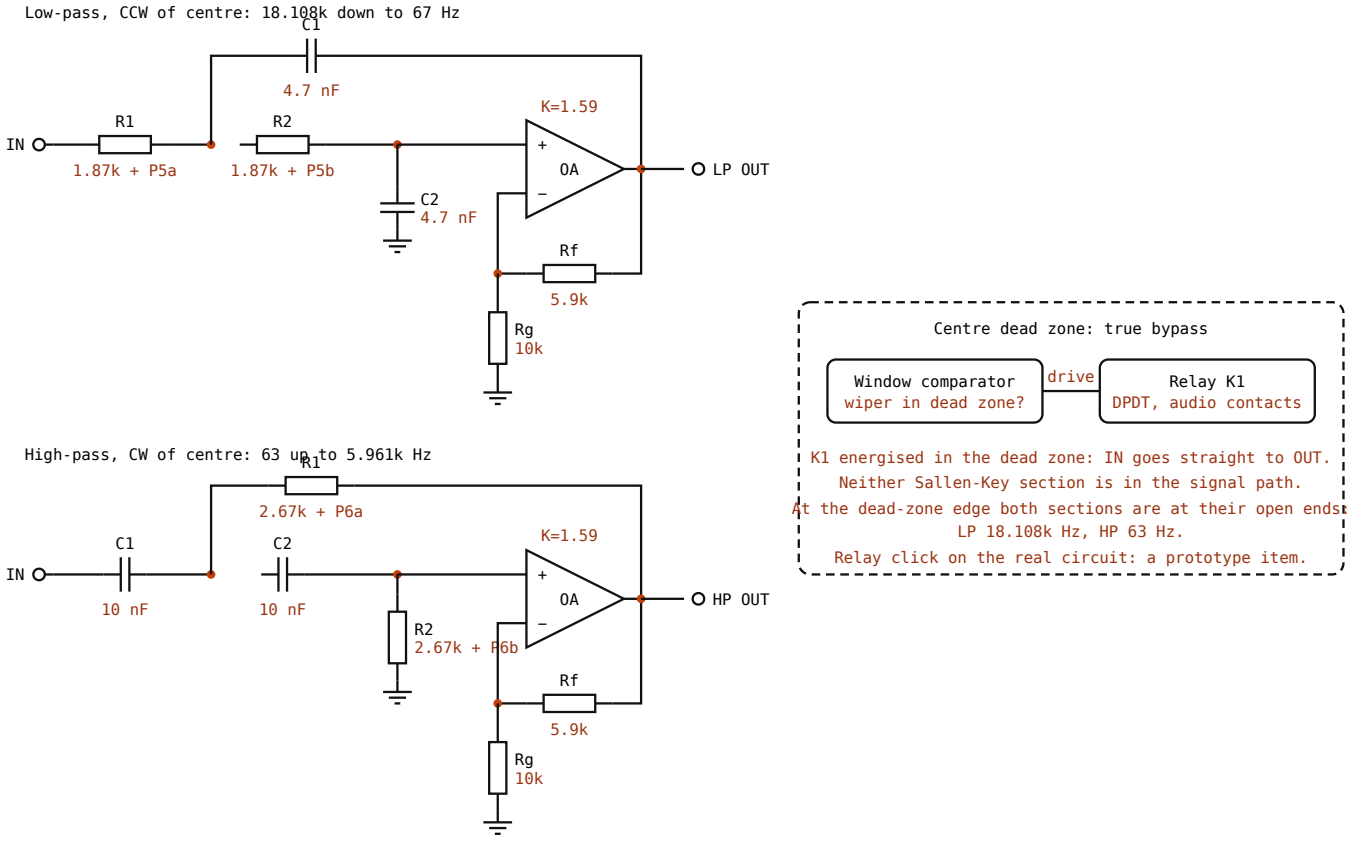
S3b THREE-BAND ISOLATOR — master section (the channel EQ reuses it at 180 Hz, see S13)



Linkwitz-Riley 24 dB/oct = two cascaded 2nd-order Butterworth (Q 0.707). Pot: min = kill, centre = unity (the x2 band amp), max = +6 dB. Allpass delay trim between bands is a prototype task.

S3b Three-band isolator (master section)

S4 CHANNEL FILTER — one-knob split filter, 12 dB per octave, relay true bypass

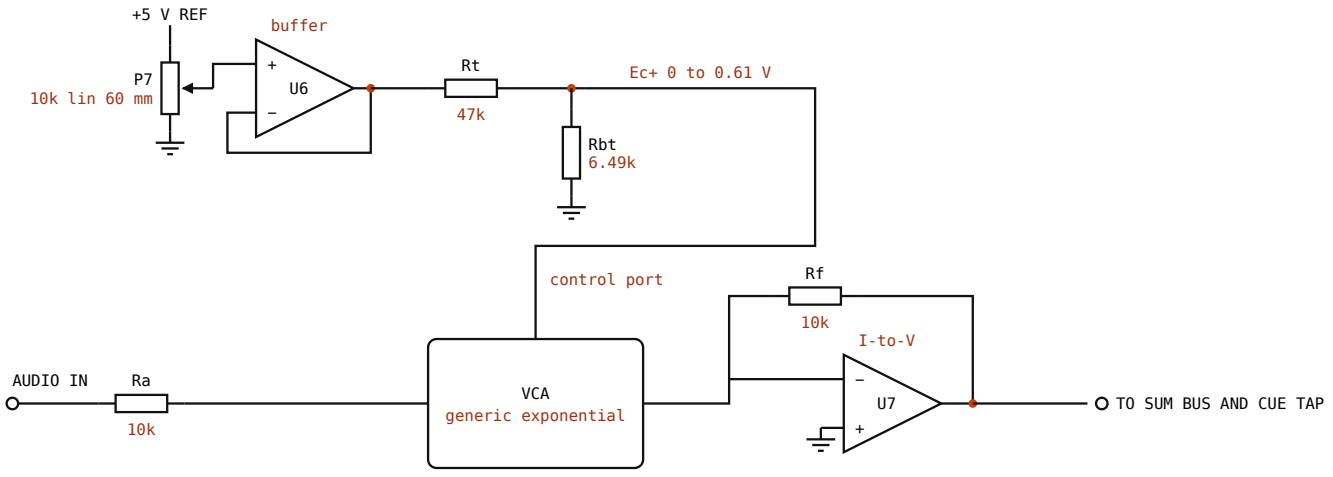


One knob, one centre dead zone. CCW of the dead zone: P5 (500k log dual-gang) sweeps the low-pass. CW: P6 (250k log dual-gang) sweeps the high-pass. fc = 1/(2πRC) and Q = 0.71 (Butterworth, low resonance, no peak).

The knob also drives the comparator through a third wiper track. The browser model on /mixer uses its own frequency range and a smoothed software bypass; it is not this circuit.

S4 Channel filter

S5 CHANNEL FADER — control-voltage VCA (no crossfader fitted)

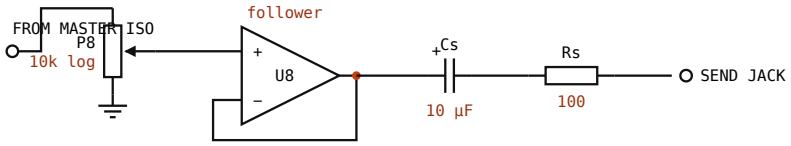


VCA control constant 6.1 mV/dB (THAT 2180-series datasheet): 100 dB range needs 0.61 V; scaler 47k / 6.49k gives 0.607 V from the 5 V reference. Rin = Rf = 10k sets 0 dB at full travel. Pin-out and bias come from the chosen VCA datasheet.

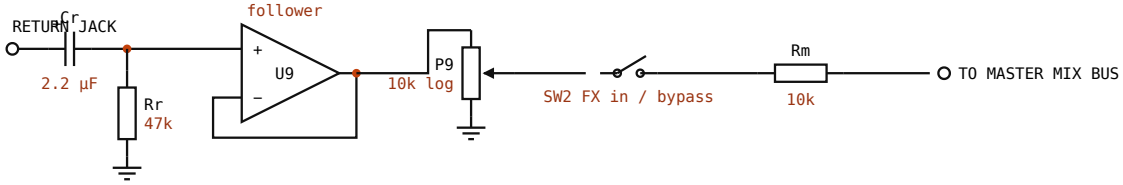
S5 Channel fader (control-voltage VCA)

S6 FX SEND / RETURN — switchable, hard-bypassed when off

Send



Return

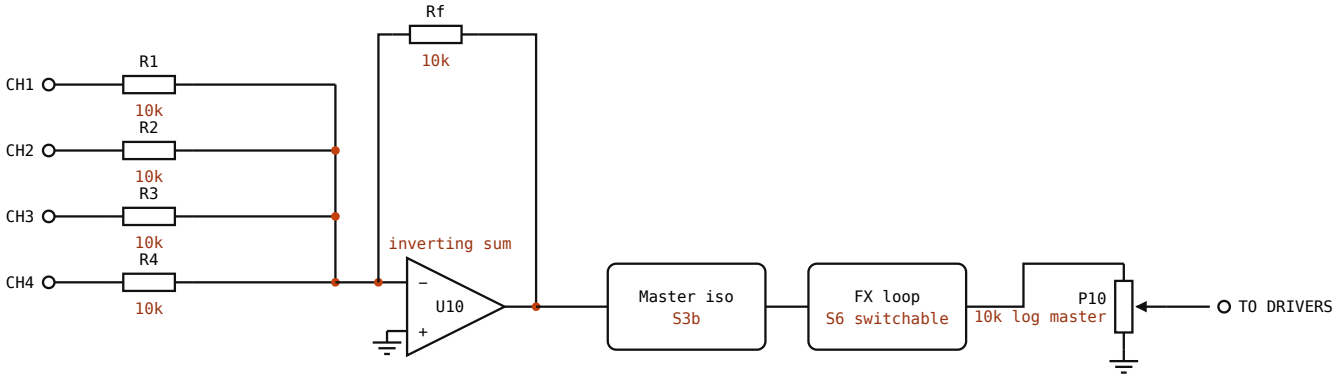


Send and return are line-level, unbalanced, nominal 0 dBV. Followers: dual low-noise op-amp. The loop is out of circuit when SW2 is off (no effect in the signal path).

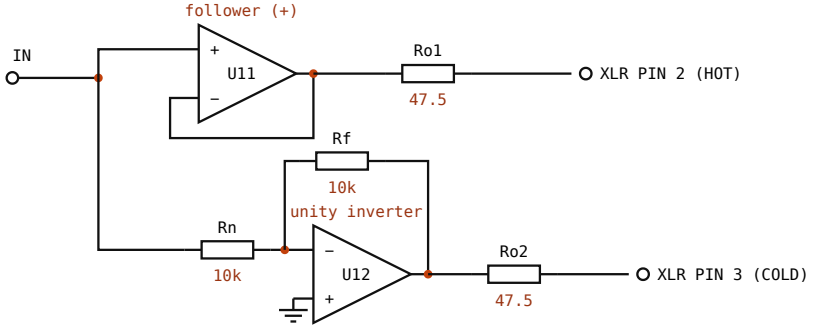
S6 FX send / return

S7 MASTER / BOOTH / CUE

Master sum and chain



Balanced driver (master XLR and booth XLR; booth has its own level pot)



Rout = 2 × 47.5 Ω = 95 Ω (impedance-balanced)

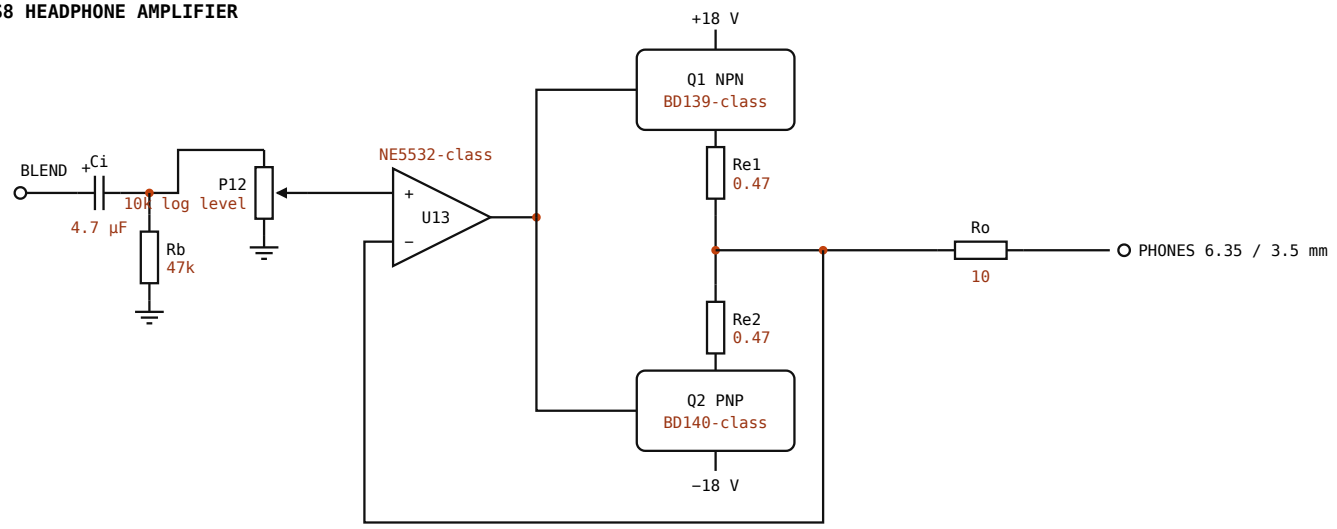
XLR pin 1: chassis via 10 Ω || 10 nF

Cue: channel pre-fader taps into an inverting sum (10k each), then cue / master blend pot 10k, then S8

All pots 10k, all sums unity per input. The sum node is a virtual earth, so crosstalk through the bus is low by construction; verify on a prototype.

S7 Master, booth and cue

S8 HEADPHONE AMPLIFIER



Calculated drive (P = V²/R):

250 mW into 32 Ω: 2.83 Vrms (4.00 Vpk), 88.4 mA rms

250 mW into 64 Ω: 4.00 Vrms (5.66 Vpk), 62.5 mA rms

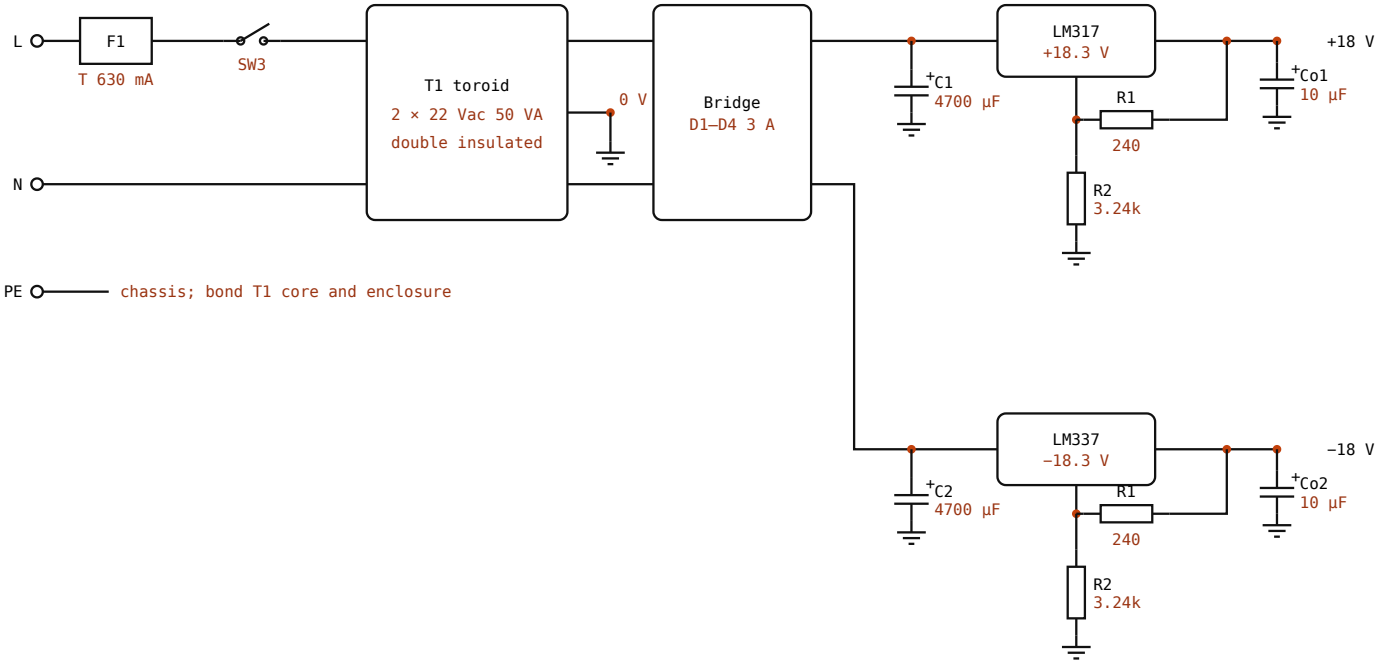
100 mW into 600 Ω: 7.75 Vrms (10.95 Vpk), 12.9 mA rms

Feedback is taken before Ro. Output-stage idle current is set at the prototype (Vbe-multiplier bias), then stabilised by the emitter resistors.

S8 Headphone amplifier

S9 POWER SUPPLY — ±18 V linear

MAINS: HAZARDOUS VOLTAGE. Build only to a qualified engineer's design review and the local safety standard.



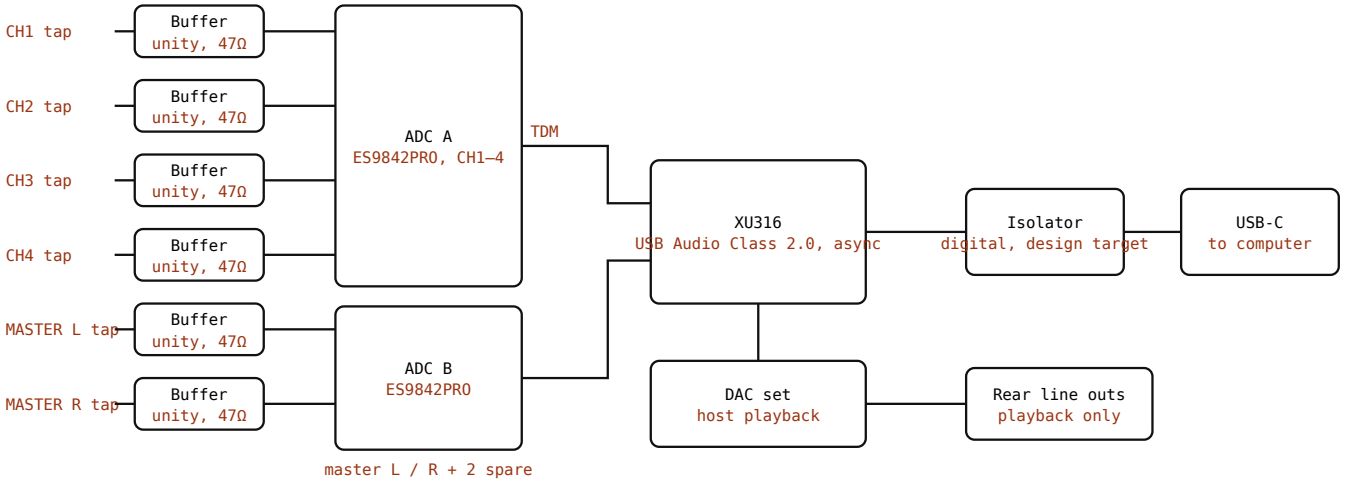
Vout = 1.25 V × (1 + R2/R1) + Iadj·R2 = 18.29 V (LM317 datasheet, Rev Z); the LM337 mirrors it with reversed diodes and capacitors.

Ripple = I/(f·C) = 1.06 V pp at 0.5 A, 100 Hz, 4700 µF. Input 29.7 V DC nominal, 28.6 V at the trough, 25.5 V at −10% mains: headroom over Vout 10.4 V (7.3 V low line); each regulator dissipates about 5.7 W: heatsink.

Also: 5.0 V precision shunt reference for the VCA control scaler; every op-amp pair decoupled with 100 nF + 10 µF per rail; star ground at the supply.

S9 Power supply

S10 USB-C INTERFACE — optional record and playback tap, design target



What this is, and is not

The analogue strips, isolators, filters and master run as in S1–S9 with the board removed or switched off: the interface only listens.

ADC: 4-channel PCM, dynamic range +122 dB in 4-channel mode (+128 dB mono), THD+N -116 dB (datasheet figures, not measured here).

122 dB ≈ 20.3 effective bits. 32-bit float is a file and transport format: its 24-bit significand covers that range without rounding.

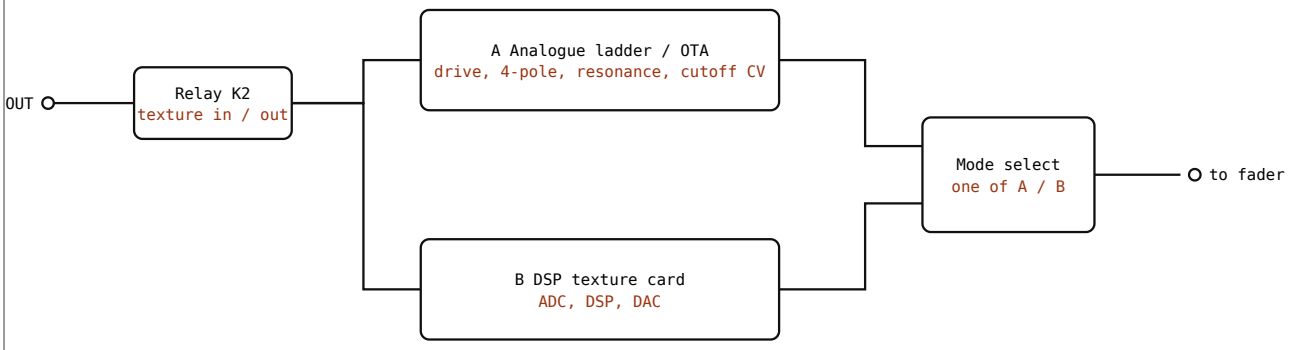
The converters deliver integer words; the host converts them to 32-bit float when it writes the file.

Raw stream at 96 kHz, 8 channels, 32-bit words: 24.6 Mbit/s (49.2 at 192 kHz); USB 2.0 high speed carries 480 Mbit/s.

Tap point per channel: before the fader (default) or after it, by a jumper. Own audio clock; USB is asynchronous.
The real range is set by the analogue noise floor and the ADC, not by the float format.
Design target. Parts named as reference choices from their datasheets; nothing here has been built or measured.

S10 USB-C interface (optional tap)

S11 TEXTURE — analogue ladder / OTA option and DSP card, design targets



When texture is off K2 routes S4 OUT directly to the fader: the path is the analogue strip of S1–S9 and no converter is in it.

A. Four-pole ladder: Stilson and Smith analyse the Moog voltage-controlled filter (24 dB per octave, resonance by feedback); Huovilainen adds the non-linear model.

B. DSP card: modes Clean SVF, Ladder (zero-delay-feedback form, after Zavalishin), Comb and Bitcrush, 4× oversampled, with drive before the filter,
resonance, and LFO or envelope-follower modulation of the cutoff. The browser model on /mixer runs this signal flow in an AudioWorklet.

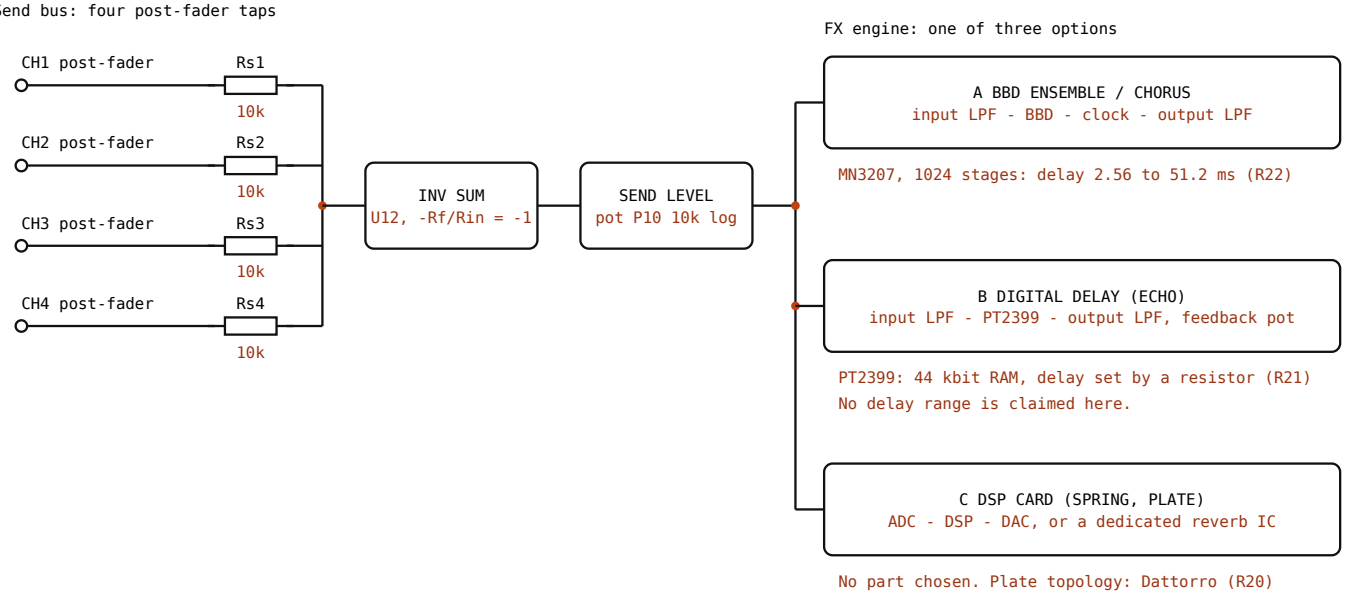
Controls: TEXTURE knob (amount; zero is off), MODE button with LED, and per-channel assign in the full design. Formant and Noise / air modes are not part of this revision.

Design target. Nothing here has been built or measured.

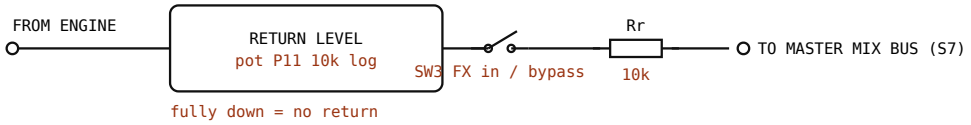
S11 Texture (optional)

S12 SPATIAL FX SEND / RETURN — optional; DESIGN TARGET, not built, not measured

Send bus: four post-fader taps



Return

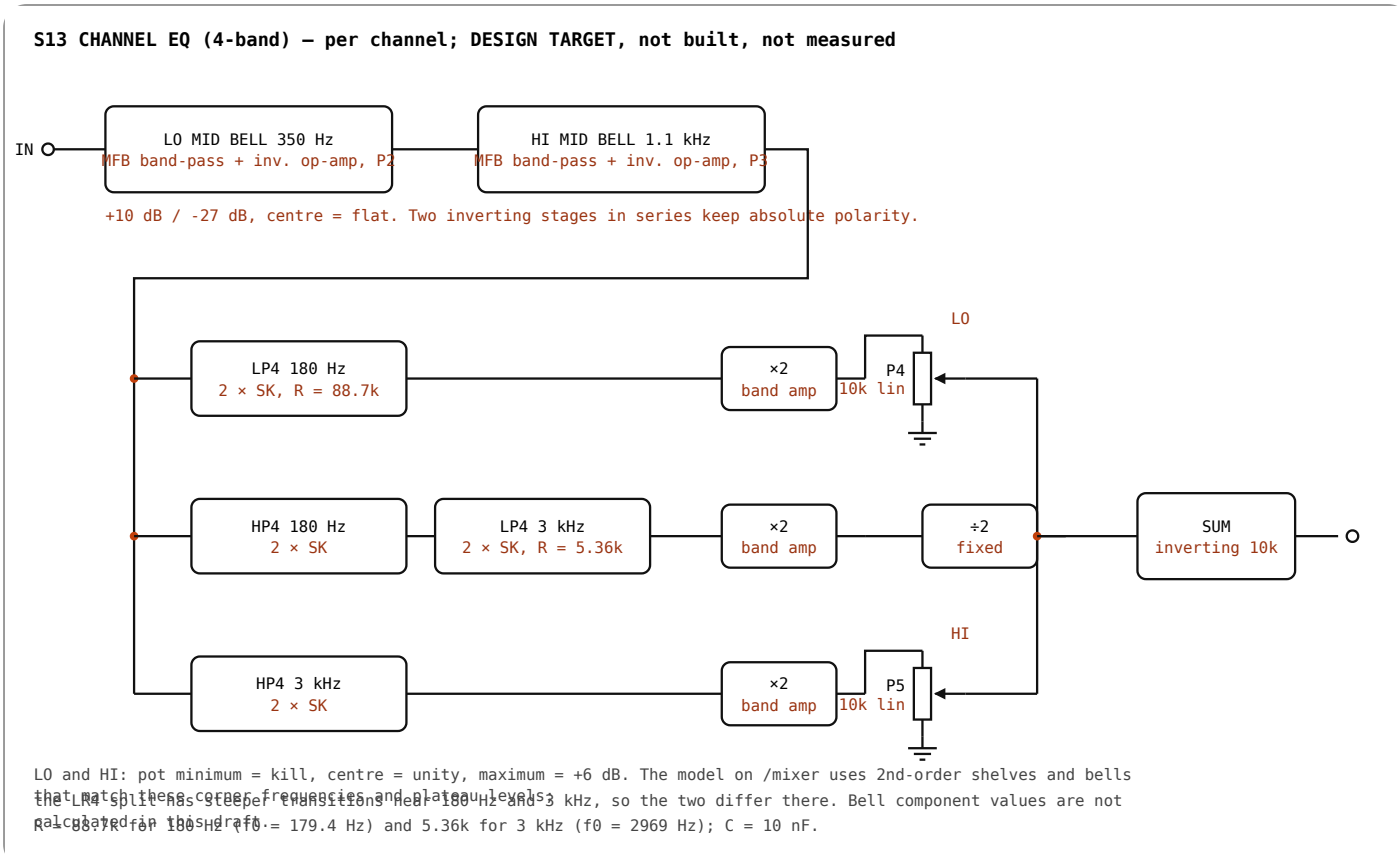


How the browser model maps to this:

Per-channel send, post-fader and post-pan, summed to a stereo bus; the return joins the mix bus before the limiter.
Return at 0 adds exact zeros in the model. In hardware SW3 removes the engine from the circuit.
The model runs four modes in floating point: tape echo, spring, plate (Dattorro) and BBD ensemble.
Hardware options A and B cover only chorus / ensemble and echo. Spring and plate need option C.
The model is not the analogue design.

BBD parts run from +5 V (R22: +4 to +10 V) and the mixer from +/-18 V: a regulator and level interface are needed.
Typical datasheet figures, each maker's own conditions: MN3207 THD 0.4 %, S/N 73 dB; PT2399 THD 0.5 %, SNR > 90 dB.
Complete circuits are not drawn.

S12 Spatial FX send / return (optional)



S13 Channel EQ (four-band, per channel)

C DESIGN RATIONALE AND DVS1 SOURCES

Each tailored feature is mapped to something he has said in public, with the source. Where the design departs from what he has said, or where nothing is documented, the table says so. The interviews are cited as the design brief, not as an endorsement, and DVS1 has not reviewed or approved the M4.

FEATURE	BECAUSE HE SAYS	THE M4 DOES	BASIS
FOUR IDENTICAL CHANNELS	*“Two Pioneer CDJs and two Technics SL-1200 turntables”* DJ Times *“2 Technics, 2 CD-Js and a mixer”* Tokafi *“adding them as a layer on a third deck, I see nothing wrong with that”* Tokafi	Four identical strips, so two decks of each kind and a third layer all have a channel.	DOCUMENTED

FEATURE	BECAUSE HE SAYS	THE M4 DOES	BASIS
PHONO OR LINE ON EVERY STRIP	*“It’s 50/50.”* DJ Times *“probably playing 50 percent off USB sticks and 50 percent off vinyl”* Little White Earbuds	A phono / line switch on each strip, because he splits his sets between records and digital files.	DOCUMENTED
A FILTER ON EVERY STRIP	*“The filters are nice.”* DJ Times	One high-pass / low-pass filter per channel, bypassed at the centre detent. He calls filters nice; he does not say he wants one per channel.	PARTLY DOCUMENTED
NOTHING BETWEEN THE DECKS AND THE SPEAKERS	*“I don’t need effects—I don’t need to fluff my mixes.”* DJ Times *“I don’t think a laptop belongs in a DJ booth.”* DJ Times	No effects in the main signal path unless you switch them in (the FX loop and the optional spatial FX are off and hard-bypassed by default), no screen, no menus, no DJ software. The optional USB-C interface is a tap off the analogue path for recording; it is not in the audio path.	DOCUMENTED
BUILT TO BE WORKED, NOT LEFT	*“I’m not letting the music play, I’m playing the music.”* DJ Mag *“you’ll never see my hands free”* DJ Mag *“I grew up on the school of DJs that change, manipulate and sculpt as they go”* DJ Mag	Large knobs and 60 mm faders, in the same place on every strip, to be ridden all set. His statements do not mention knob size or spacing; that part is the designer’s inference.	PARTLY DOCUMENTED
CHANNEL FADERS	*“manipulating things by volume, EQ, fader levels”* XLR8R *“until he puts that fader up”* DJ Times	A linear channel fader on each strip. He talks about volume and fader levels; he does not say whether he prefers linear or rotary faders.	PARTLY DOCUMENTED

FEATURE	BECAUSE HE SAYS	THE M4 DOES	BASIS
EQ TO FIT A THIRD TRACK IN	*"Because I have at least two levels going at any time and am looking for a third track to add in, I'm always pushed to manipulate or EQ things to make them fit together."* XLR8R *"He sometimes ducked a high-mid on the 3s and 4s before bringing it back in for the 1s and 2s, something that the Xone:96 mixer he was using allowed him to do."* Magical Dancefloors (writer's account)	Shelves that go to full kill (and bells that cut 27 dB), so a band can be cleared for the next track. He says he EQs to make tracks fit; how far he cuts is not documented. The second quote is a writer's account of one show.	PARTLY DOCUMENTED
HEAR THE ROOM IN THE BOOTH	*"I need subs, and tops in the booth to really feel that."* Tokafi *"the monitors in the booth are so in tune with the dance floor sound that you know you can be really precise"* XLR8R *"Monitor Volume must be accessible to talent during the performance"* Rider 2019	A BOOTH level knob on the panel and a separate balanced booth output, so the monitors are controlled from the mixer. The M4 supplies no speakers.	DOCUMENTED
NO COMPRESSOR IN THE MONITOR PATH	*"NO COMPRESSOR ON MONITORS."* Rider 2019	No compressor or limiter anywhere in the master or booth path of the hardware design. (The browser demonstration has an output limiter; see the context list.)	PARTLY DOCUMENTED
DARK BOOTH	*"NO Lights on DJ, Artist prefers to be in the dark"* Rider 2016	A METER DIM switch that turns the LEDs and meter lamps down to a low amber.	PARTLY DOCUMENTED

FEATURE	BECAUSE HE SAYS	THE M4 DOES	BASIS
HEADPHONES	*“the Sennheiser HD 25-2s, for headphones”* Little White Earbuds	A PHONES output on 6.35 mm and 3.5 mm, sized for DJ headphones of the kind he names; both outputs can be used together, with split cue. How he cues is not documented, so cue / mix, split cue and the dual outputs are the designer’s choice.	PARTLY DOCUMENTED
VERY LONG SETS	*“the ability to spread out and play eight- or 10-hour sets”* DJ Times	External linear supply, no firmware and no boot sequence, user-replaceable faders and pots. His statement is about the length of his sets; these are inferences about what suits them.	PARTLY DOCUMENTED

Where the design departs, or nothing is documented

- **FX loop.** He says “I don’t need effects” (DJ Times, 2014), yet his own riders ask for stereo cables for an effects box that he brings, on aux 1 (2016 rider, q17). The M4 therefore has an FX send and return, hard-bypassed and off until it is switched in. This departs from the 2014 statement and follows the riders.
- **USB-C interface.** He says “I don’t think a laptop belongs in a DJ booth” (q7). The optional USB-C record / playback interface departs from that: it is a tap for recording, off by default, with no DJ software and no screen, and nothing in his documented setup asks for it.
- **Spatial FX.** He says “I don’t need effects—I don’t need to fluff my mixes” (q4) and his riders ask only for a plain effects send and return that he fills with his own box (q17). DVS1 favours minimal effects, so the M4 spatial FX (tape echo, spring, plate, BBD ensemble) are the designer’s addition, not something he asked for. They are optional, off with the return at zero, and hard-bypassed in hardware by SW3. No source reviewed mentions spring, plate or tape echo.
- **Four-band channel EQ.** He is quoted as saying “I don’t need a 4-band EQ” (q6). The M4 channel EQ has four bands anyway (HI, HI MID, LO MID, LO), which departs from that statement; it was added at the site owner’s request. The only related account is a writer reporting him ducking a high-mid on an Xone:96 (q26). The frequencies and ranges follow the published Xone:96 specification, not anything DVS1 said.
- **Texture section.** No source reviewed mentions resonant, ladder, comb or bitcrush effects, and he says “I don’t need effects” (q4). The texture section is the designer’s addition; it is off by default and exactly bypassed when off.
- **Four phono stages.** The documented setup has two turntables (q1, q2), and the 2019 festival rider lists none. Four phono stages go beyond what he has said he needs; they are there so any deck can use any channel, and could be cut to two to save cost.
- **Layout.** He says he does not like the layout of the Xone mixers and does not say what he dislikes (q3). The M4 strip layout is the designer’s guess. Nothing documented says it is what he wants.
- **Crossfader.** None of the sources reviewed says whether he uses a crossfader. The M4 has none; that is a design choice, not something he asked for.
- **Rotary or linear faders.** No statement about rotary-fader mixers was found. The M4 uses linear faders (see above).
- **Master isolator and VU meters.** No source reviewed mentions either. A writer reports him ducking a high-mid on an Xone:96 (q26); that is the only related account.
- **Output level and headroom.** The +28 dBu maximum output and the ±18 V rails are an engineering choice. He has not said anything about mixer headroom.
- **Pitch and tempo.** He pitches records on Technics MK5G decks (q24). Pitch is on the turntable, so the M4 hardware has no tempo, sync or BPM function. The browser app adds BPM detection, tap tempo, a pitch fader, loops, hot cues and an

optional sync as tools of the software model; they are a departure from this, not something he asked for, and sync is off until pressed.

- **His riders name a mixer.** Both riders ask for an Allen & Heath Xone (92, then 96) and say “STRICTLY NO OTHERS” (q30). The M4 would not meet his rider, and nothing here suggests he would use it.
- **Not found.** No usable interview text was found from Resident Advisor, Mixmag or Pioneer on mixer layout, rotary mixers, headphone cueing or booth monitoring. A 2016 Slam Academy video interview exists but was not transcribed, so nothing from it is used.

Context

- **Browser demonstration.** The Web Audio demonstration on this page has an output limiter and soft saturation so that nothing clips in a browser. The hardware design has no compressor or limiter in the master or booth path.
- **Not applicable.** His remark about keeping “a maximum of seven channels” in DJ Mag (2018) is about his studio, not his DJ mixer, and is not used.
- **Dates.** The sources run from 2013 to 2019. His setup has changed over that time (for example Xone 92 to 96, CDJ 2000 to 3000), so these are statements from their date, not today.

Sources

1. **[DT]** Jim Tremayne, “DVS1: Back to the Old School”, DJ Times, 5 February 2014 (archived copy) — open (https://web.archive.org/web/20260315170802/https://www.djtimes.com/2014/02/back-old-school/)
2. **[DM]** Rob McCallum, “DVS1: The techno purist rages against the machine”, DJ Mag, 31 January 2018 — open (https://djmag.com/content/dvs1-techno-purist-rages-against-machine)
3. **[LWE]** Steve Mizek, “DJ Debriefing: DVS1”, Little White Earbuds, 21 August 2013 (archived copy) — open (https://web.archive.org/web/2024/http://www.littlewhiteearbuds.com/feature/dj-debriefing-dvs1/)
4. **[TK]** Tobias Fischer, “Interview with DVS1”, Tokafi (undated on the page) — open (https://www.tokafi.com/15questions/interview-dvs1/)
5. **[XR]** “Ask the Experts: DVS1”, XLR8R, 24 March 2017 — open (https://xlr8r.com/features/ask-the-experts-dvs1/)
6. **[R16]** Triangle Agency, artist tech rider, DVS1 – DJ (PDF, 2016 upload) — open (https://triangleagency.com/wp-content/uploads/2016/04/dvs1_techrider.pdf)
7. **[R19]** Triangle Agency, artist festival tech rider, DVS1 (DJ) (PDF, 2019 upload) — open (https://triangleagency.com/wp-content/uploads/2019/04/DVS1_TechRider_Full_2019_FESTIVAL.pdf)
8. **[MD]** “DVS1 Wall of Sound dominates Los Angeles”, Magical Dancefloors. A writer’s account of a show, not DVS1’s words — open (https://www.magicaldancefloors.com/p/dvs1-wall-of-sound-dominates-los)

Quotations are word for word from the linked pages (checked by script against the page text). Everything else on this page is the designer’s own wording. They are cited as the design brief, not as an endorsement, and DVS1 has not reviewed or approved the M4.

D

BROWSER MODEL AUDIO SPECIFICATION (SIMULATION)

These are simulation results for the browser model, not measurements of the analogue mixer and not hardware measurements. They were produced by running the same AudioWorklet code in Node at 48 kHz (Float64 inside the filters, Float32 at the module boundaries) with the script `tests/dsp_test.js`. They describe the software on /mixer; the circuits in this manual are design targets and have no measured figures.

TEST	CONDITION	RESULT (SIMULATION)
THD+N	1 kHz (999.76 Hz, exactly on an FFT bin), −6 dBFS, 20 Hz to 20 kHz, clean path	−153.8 dB. With the oversampling filters forced on: −147.7 dB. The Float32 test tone alone: −153.8 dB, so this is the limit of the measurement
Null test, bypass against input	Noise at about −18 dBFS rms, low-passed at 16 kHz, through EQ, filter, texture, both saturators and the limiter, all at rest	Output equals the input delayed 288 samples, bit for bit. With the oversampling filters forced on: residual −76.3 dB of the input
Noise floor	Digital silence in; then a −60 dBFS 1 kHz tone with the filters forced on	Silence in gives exactly 0 out on every sample. With the tone, everything else from 20 Hz to 20 kHz is at −206.3 dBFS (Float32 rounding)
Aliasing, texture drive	15 kHz sine at −6 dBFS, drive at maximum, everything from 20 Hz to 20 kHz except the tone, relative to the tone	No oversampling −12.6 dB; 2× −29.2 dB; 4× −54.6 dB
Aliasing, saturator	15 kHz sine at +6 dBFS into the knee	No oversampling −58.8 dB; 2× −113.4 dB; 4× −117.0 dB
Frequency response	All knobs at 12:00, 14 tones from 20 Hz to 20 kHz at −12 dBFS	Clean path: 0.0000 dB deviation. With the oversampling filters forced on: worst 0.006 dB at 20000 Hz
Limiter, true peak	White noise to Nyquist at −3 dBFS rms, read by an independent 16-phase true-peak meter	−1.1 dBTP against the −1.0 dBTP ceiling. The detector carries a +0.3 dB margin, so loud programme is held at about −1.3 dBTP
24-bit export	Quantisation error of the dithered export against its theory	Error 0.499 LSB rms against 0.500 expected for triangular dither; mean 0.0020 LSB; clipped samples are counted and reported to the user
CPU, desktop only	One channel in Node on the build machine, multiples of real time	Saturator 2× 157, 4× 76; texture ladder 2× 24, 4× 14. No phone was measured

What the engine does

- **Sample rate.** The audio context runs at the device's rate by default (48 kHz or 96 kHz where the device offers it) or at a rate the user picks; nothing is resampled inside the engine. The latency hint is the user's choice (interactive, balanced or playback).
- **Precision.** 32-bit float at the module boundaries. The recorder writes 32-bit float WAV at the context rate; the optional 24-bit file uses triangular (TPDF) dither and reports any samples that were over full scale.
- **Nonlinear stages.** The texture drive, the strip saturators and the master saturator are oversampled 4× (2× in the efficient mode) with half-band filters. The saturator is the identity below 0 dBFS, so below that level nothing is shaped and the bypass is bit-exact; above it the curve bends smoothly towards +18 dB of internal headroom. The spatial FX's own tape, spring and BBD saturators are **not** oversampled yet; they run at the base rate inside feedback loops.
- **Smoothing, DC and denormals.** EQ gains and filter positions are smoothed over about 12 ms; the texture wet path has a 5 Hz DC blocker; the texture filter states are flushed to zero below 1e-25 so no denormals arise.
- **Limiter.** A look-ahead true-peak limiter at −1 dBTP on the output only, using an 8× detector. It does not touch the signal below its threshold. Channel strips have no limiter.
- **CPU.** The master watches whether render callbacks keep up with real time. If they fall behind three windows in a row the engine switches itself from 4× to 2× and says so on the page. Phones and machines with four cores or fewer start at 2×.

- **Track decks.** Each of the four decks takes a lossless file from the user's own device: WAV (16, 24, 32-bit integer or 32-bit float), AIFF or AIFF-C (PCM, little-endian "sowt" or float) and FLAC. Anything else, including MP3 and other compressed formats, is refused by reading the file's header, not its name. WAV and AIFF are decoded by the page's own reader into 32-bit float without rounding (a 16 or 24-bit sample divided by 2^{bits-1} is exact in Float32); FLAC is decoded by the browser. The file stays in the browser tab: nothing is uploaded and the page makes no request that carries it. A loaded deck replaces its test signal and runs through the strip's EQ, filter, texture and saturator, the FX sends, the master and the recorder. The strip processes a stereo track with two identical chains. The deck's panner is centred for a track, so a stereo track enters the master at unity on each side, while the mono test signals sit at −3 dB per side. When a file's sample rate differs from the engine's, the browser converts it as it plays (the deck says so); there is no "no resampling" claim for decks. Rates from 8 to 96 kHz are accepted.
- **The radio.** The site's radio player uses the highest-quality stream each provider officially offers; some provider previews are lossy and the site does not re-host them.

Limits of this table: a simulation of the code, a single noise seed, ideal 32-bit and 64-bit arithmetic rather than a browser's audio stack, and no converter or analogue stage. Nothing here supports a claim about how it sounds.

E

BROWSER APP: DJ TOOLS, SESSIONS AND OFFLINE (SIMULATION)

These are checks of the browser app on synthetic test signals made for the test. They are not measurements of the analogue mixer, they use no real music, and they say nothing about how the app sounds. The app runs only in the user's browser; nothing is uploaded.

What it adds to the browser model

- **Install and offline.** A web app manifest and a service worker limited to /mixer. After one online visit the page, scripts, fonts and images of /mixer load without a network. The other pages of the site, the documents and the fan room are not cached by it. The cache is versioned and older versions are removed when the worker updates.
- **Saved sessions.** The mixer, EQ, FX, hi-fi and deck settings are kept in the browser's localStorage and restored on reload, can be saved under a name, and can be exported and imported as JSON (settings and cue points only, no audio). Tracks are stored only if the user ticks the opt-in box, then in the browser's IndexedDB, and can be removed again.
- **Per-deck tools.** Tempo detection, tap tempo, a pitch fader of ±8% or ±16%, nudge, loops of 1/4 to 16 bars on the detected beat grid, four hot cues and a sync option that is off until pressed. Pitch and tempo move together, like a turntable: key lock (changing tempo without changing pitch) is **not available** in this version.
- **Library, MIDI, keys.** A searchable list of the tracks loaded in the tab. Web MIDI learn and map with save and load, where the browser has Web MIDI (not Safari). Keyboard shortcuts with a help sheet, touch gestures, and a first-run tour that can be skipped.

The hardware design is unchanged. Appendix C records that DVS1 pitches records on the turntable (q24) and that the M4 hardware has no tempo, BPM or sync function. The app's tempo tools are an addition of the software model, not something DVS1 asked for, and the sync option is off by default.

Tempo detection on synthetic tracks

TRACK (GENERATED)	TRUE BPM	READ BPM	ERROR (BPM)	FIRST BEAT ERROR (MS)
click 100	100.0	100.000	0.000	0.250
click 128	128.0	128.000	0.000	0.270

TRACK (GENERATED)	TRUE BPM	READ BPM	ERROR (BPM)	FIRST BEAT ERROR (MS)
click 140	140.0	140.000	0.000	0.280
kick 174	174.0	174.000	0.000	0.680
kick 124	124.0	124.000	0.000	0.660

Run in Chromium through the real app. In addition, 80 synthetic tracks (click and kick plus hat, 80 to 180 BPM, 44.1 and 48 kHz, 40 s) were run through the same detector in Node: 78 read within 0.05 BPM, worst error 0.01 BPM, worst first-beat error 1.67 ms. The misses are octave errors: kick 80 read as 160. The detector looks at the first 150 s of a track, searches 70 to 190 BPM and assumes a steady tempo; real recordings with swing, tempo drift or no clear kick were not tested.

Loops, hot cues, pitch and sync

CHECK	CONDITION	RESULT
Loop sample accuracy	A 2.0 s loop on a ramp track, rendered offline at rate 1.0, three start positions	Worst error 0.007 samples over 265 377 samples
Hot cue jump	Cue stored as a whole sample, jump rendered offline	Worst error 0.013 samples; stored frame equals the frame at the press
Pitch +5%	Position read back after 20 000 output samples, offline	Worst error 0.007 samples
Tap tempo	Seven taps 500 ms apart, driven by a page timer	119.37 BPM (120 expected); the error is the timer’s own jitter
Sync	Two playing decks, follower tempo matched to the leader and its beat phase aligned when sync is pressed; measured against the detected beat grids, so it is only as good as the detection	Tempo difference 0.0 BPM, beat offset -0.0 ms at the press, unchanged 2.5 s later
Interaction speed	Four decks playing, input to the next frame callback, 80 samples, headless Chromium with software graphics (a frame takes about 40 ms there)	Median 40 ms, 95th percentile 46 ms, worst 88 ms. The handlers themselves took 1.4 ms at the 95th percentile and never more than 4.4 ms. The browser’s own event timing, which waits for the next paint, reported up to 152 ms in this setup; in a side-by-side run of thirty clicks and slider moves with four decks playing, this page and the previous release both reported 104 ms at most

Limits

- Live playback at a pitch other than 0% is positioned by the browser’s resampler; the sample-exact figures above are for rate 1.0 (loops, cues) and for the offline position check at +5%.
- The loop and cue figures are for an offline render of the same source node settings the decks use; the audio device and the browser’s scheduling add their own latency on a live deck, which was not measured.
- MIDI learn takes absolute controllers (knobs and pads). Relative encoders, pickup and 14-bit controls are not supported. It was tested with a scripted controller, not hardware.
- Only Chromium was tested. Safari has no Web MIDI; other browsers were not run.
- Interaction timing was measured on a build machine without a graphics card, so it says little about phones.