

SAK-TUBE

*An amplifier.
A connected circuit.
A little more life.*



REFERENCE MANUAL

Version 0.1.3 · Windows x64 · VST3 & standalone

Controls, circuit architecture, parameter provenance and model limitations.

Read this first

SAK-TUBE is a valve-amplifier effect built around an interacting electrical circuit. Two loaded triode stages feed a long-tailed phase inverter, a push-pull output pair, an output transformer and a reactive speaker load. Their supply, coupling and magnetic states persist while you play.

The current release always uses the Jiles–Atherton output core, the Musical Drive/Power structure and parameter-only level compensation. These are the product architecture, not optional historical modes. This manual describes **version 0.1.3**, including the revised high-frequency response.

Find your way

Page	Topic
3	Installation, removal and first session
4	Main controls and practical settings
5	Valve choices and Circuit Details
6	Complete signal flow and gain structure
7	Input transformer and loaded preamp
8	Tube-current laws and their evidence
9	Tone, inverter, grid conduction and feedback
10	Output valves and transformer equivalent
11	Jiles–Atherton magnetic state
12	Reactive load, supply and memory
13	Frequency response and sample rate
14	Validation results and known limitations
15	Reference and provenance index
16	Open source, rebuilding and troubleshooting

How to interpret the circuit

Reference-backed means a value or curve is available in a named source. **Fitted** means parameters were identified against selected reference data. **Engineering choice / approximation** means a circuit choice or reduction is plausible but is not a measurement of the particular hardware being represented. A fitted component is not proof that the whole amplifier duplicates a real physical unit.

SAK-TUBE is an original amplifier configuration. Manufacturer and tube-family names identify technical references; they do not imply endorsement. Internal voltages and currents are simulated SI quantities, not hazardous physical outputs.

Electrical load, not cabinet audio. The speaker impedance affects the power tubes, transformer and feedback. No acoustic cabinet, microphone, room response or convolution IR is included. For a recorded guitar-cabinet sound, add your preferred cabinet IR after SAK-TUBE.

Install. Open. Play.

Windows installation

1. Close your audio host. Run **SAK-TUBE-0.1.3-Setup.exe**. The release is unsigned, so Windows may show an unknown-publisher warning. Obtain it from your trusted SAK Audio download and compare the published SHA-256 checksum if you want to verify the file.
2. The installer places the complete VST3 bundle in **C:\Program Files\Common Files\VST3\SAK-TUBE.vst3** by default. Administrator access is required for this shared location. You can choose another VST3 folder during setup.
3. The standalone application is optional. The reference manual, license and notices are installed with the application support files. Setup installs the Microsoft Visual C++ x64 runtime if the required version is missing; it does not remove that shared runtime on uninstall.
4. Start your host and rescan VST3 effects if needed. Insert **SAK-TUBE** on an audio track. The supported layouts are mono-to-mono and stereo-to-stereo.

Tested platform: Windows 11 x64. The installer accepts Windows 10 or later; Windows 10 has not been independently validated for this release. A host with 64-bit VST3 support is required for the plugin. This package does not contain macOS, Linux, AU, VST2 or AAX builds.

Portable installation and removal

Extract the ZIP, then copy the **whole SAK-TUBE.vst3 folder**, including its Contents directory, to your VST3 location. Do not copy only the inner binary. If the host reports a missing VCRUNTIME/MSVCP library, install the official Microsoft Visual C++ x64 Redistributable linked in the package's installation notes.

For an installer installation, use Windows **Settings** → **Apps** → **SAK-TUBE** → **Uninstall**. For a manual installation, close the host and remove the bundle you copied. Projects and host-owned preset files are not removed by the installer.

A useful first session

Start with the defaults: **12AX7 / EL34 / GZ34**, Drive **3.5**, Power **5.0**, Tone **5.0**, Presence **4.0**, Sag **5.0**, Output **-12 dB**, Quality **1***. Feed a normal instrument recording or track at a sensible level. Source level matters: a mastered mix excites the circuit much harder than a quiet guitar take.

Lower Drive for less preamp excitation. Raise Power to hear the output section work. Use Output for listening level. The compensation reduces level swings; it does not guarantee constant loudness or prevent digital clipping. If the red **CLIP** indicator lights, lower Output.

Double-click a knob to restore its default. Hold Shift while dragging for fine adjustment. The standalone application's audio-device settings select your interface and buffer size. Avoid routing its output back to its input.

The six main controls

Control	Default	What changes
Drive	3.5 / 10	Source excitation before the input transformer and preamp. Valve-dependent finite taper; makeup follows the preamp.
Power	5.0 / 10	The loaded interstage control feeding the inverter/output section. Makeup follows the entire physical output circuit.
Tone	5.0 / 10	A passive treble cut before the inverter. Left is dark; right progressively opens the bandwidth.
Presence	4.0 / 10	The high-frequency character of the negative-feedback path. Its effect depends on Feedback and the driven circuit.
Sag	5.0 / 10	Effective rectifier/source resistance: left is firmer, right gives more supply movement.
Output	−12 dB	Final monitoring gain, after the speaker-terminal simulation. Range −36 to +6 dB.

Drive and Power at zero are **finite clean settings, not mute**. Use the host's bypass or track mute to silence the effect. The Output meter shows monitoring level; the B+ readout follows the simulated shared reservoir rail.

Starting points, not fixed promises

These positions use the visible 0–10 scales. Keep Tone around 5, begin with the default valves, then adjust to your source. A strong transient can cross several stage thresholds even at a modest knob position.

Intention	Drive	Power	Listen for
Light coloration	1–3	1–3	Attack retained, small harmonic changes.
Preamp edge	4–6	1–3	Increasing hair while the output section stays relatively restrained.
Output-stage bloom	1–3	7–9	More output compression, transformer excitation and recovery.
Dense crunch	6–8	4–6	Preamp saturation with moderate output movement.
Full overload	9–10	9–10	Both sections heavily excited; strong blocking and supply interactions are possible.

On a bass part, compare low Drive/high Power against the reverse. Low notes demand more flux linkage for the same primary voltage, making transformer stress and supply recovery more apparent. On a mix or drum bus, start with low Drive and Power; SAK-TUBE has no wet/dry control, so use a host parallel bus if desired and let the host compensate its reported latency.

Tone is not a post-processing brightness shelf. At the centre its modeled RC corner is approximately **41.1 kHz**. At the dark end it reaches **650 Hz**; near the bright end that particular cut is effectively open. The other physical networks still impose their own response.

Valves and Circuit Details

The valve selectors choose coherent amplifier configurations. They are not a claim that these devices can safely be swapped into one physical amplifier without circuit changes.

12AX7 has the highest intrinsic gain and the easiest access to preamp saturation. **12AT7** uses its own fitted current law, capacitances and bias/loading, with additional available input gain. **12AU7** retains a much lower intrinsic gain and greater input headroom; its wider source-gain range makes later-stage drive accessible. Makeup does not replace their electrical differences.

EL34, 6L6GC, EL84 and KT88 select different plate/screen laws and matched supply, bias, screen resistance and primary load. The phase inverter uses the selected triode family. **Silicon, GZ34 and 5U4G** select progressively larger nominal source drop and resistance; these are cycle-averaged rectifier equivalents, not individual rectifier plate-curve fits.



Detail	Default / range	Meaning
Bias	5.0 / 0–10	Changes negative-bias target. Higher settings reduce its magnitude, generally increasing idle conduction. Supply and thermal state still affect it.
Feedback	3.5 / 0–10	Negative feedback around the output section; changes damping, gain and distortion interaction.
Iron	5.0 / 0–10	Changes transformer saturation strength/geometry equivalent. It is not a measured steel-grade selector.
Reservoir	47 μ F / 16–150 μ F	Main supply storage per channel. More storage changes droop and recovery; it does not eliminate source resistance.
Coupling	22 nF / 4.7–100 nF	Preamp coupling capacitance. Alters low-frequency transfer and charge memory under grid current.
Quality	1 $\times$ / 1 $\times$, 2 $\times$, 4 $\times$	Circuit execution and alias rejection; 1 $\times$ uses two full steps, 2 $\times$ splits preamp/output rates, and 4 $\times$ retains the reference.

Valve, rectifier and quality changes briefly fade the monitoring signal to avoid a discontinuity. Ordinary control changes are smoothed. No separate Legacy, simplified-core or raw-level product mode is exposed.

One connected signal path

INPUT → finite Drive gain → INPUT TRANSFORMER → PREAMP 1 → loaded coupling → PREAMP 2 → Power divider / coupling → TONE → Drive makeup → PHASE INVERTER → power-grid coupling → OUTPUT VALVES ↔ OUTPUT TRANSFORMER ↔ REACTIVE SPEAKER → Power makeup → OUTPUT

Negative feedback returns from the electrical output toward the inverter drive. A four-node supply serves output plates, screens, inverter and preamp. Both stereo channels draw from the same supply, with storage and source scaling for channel count.

Drive mapping

Input conversion is modeled $\text{source} = \text{input} \times 0.18 \times 10^{(\text{Drive_dB}/20)}$. The 0.18 factor is a digital-to-circuit calibration choice, not a universal relationship between dBFS and guitar pickup voltage. The following electrical gain knots use normalized knob position; the UI displays that position multiplied by ten.

Position	0	0.25	0.50	0.75	1.00
12AX7 gain, dB	-18	-14	-6	12	42
12AT7 gain, dB	-18	-14	-6	14	48
12AU7 gain, dB	-6	-2	6	28	54

The finite low end leaves useful signal for compensation; no setting approaches negative-infinite input gain. The taper gives substantial travel to clean and edge operation while retaining overload at the top.

Power and two makeup locations

For normalized Power u , the circuit position is $p = 0.25 + 0.75u$, with wiper $w = p^3$. The equivalent series branch is $22 \text{ k}\Omega + 1 \text{ M}\Omega(1 - w)$ and its leak is $470 \text{ k}\Omega \parallel (1 \text{ M}\Omega \times w)$. This changes excitation and loading before the nonlinear output section.

Drive makeup is applied **after the loaded preamp coupling/Tone network and before inverter/NFB drive**. Power makeup is applied **after the full electrical output transformer and reactive load**, outside the feedback loop. Output is the final user trim.

The compensation tables were calibrated from circuit measurements, including 200 Hz reference measurements, for all 12 triode/output-family combinations. Drive uses 12 knots per combination. Power uses a 5-Drive $\times$ 9-Power grid per combination. Gains interpolate in dB, are bounded to **-40...+24 dB**, and smooth over approximately **50 ms**. There is no envelope follower, AGC or signal-dependent gain reduction.

These tables aim for useful approximate levels. They do not match arbitrary input spectra, transients or perceived loudness exactly. A signal-domain makeup stage is also an engineering partition: its reverse impedance and upstream grid-current loading are not a fully bidirectional physical amplifier stage.

Input iron and loaded triodes

Input-transformer equivalent

The 1:1 input equivalent uses Jensen JT-11P-1 data [1]. Published primary/secondary DCRs are **1.45 / 1.55 kΩ**, with case capacitances **98 / 110 pF**. A **600 Ω** source and the modeled input grid network complete the circuit. The **300 H** magnetizing and **12 mH** leakage inductances are inferred; the **208 pF** shunt is an engineering lump of the published case capacitances, not an exact multiwinding capacitance matrix.

The input core retains a bounded cubic/play hysteretic reduction, with nominal knee linkage **0.08 V·s** and play width **0.0003 V·s**. It is distinct from the Jiles–Atherton output core. The input grid has a **33 kΩ** stopper and **1 MΩ** leak. These networks establish source loading before nonlinear amplification.

Preamp configurations

Plate and cathode resistances are amplifier design choices. Brimar RC-coupled examples and published gain-stage analysis informed the alternatives [8]; the table does not redefine intrinsic tube properties. Each stage has a **22 μF** cathode bypass capacitor.

Valve	Plate R, stages 1 / 2	Cathode R, stages 1 / 2	Interstage series / leak
12AX7	100 / 82 kΩ	1.5 / 2.2 kΩ	680 / 150 kΩ
12AT7	100 / 100 kΩ	3.3 / 3.3 kΩ	470 / 330 kΩ
12AU7	100 / 100 kΩ	4.3 / 4.3 kΩ	330 / 470 kΩ

Coupling capacitors hold actual charge. Source impedance, leak resistance and grid conduction determine both the small-signal corner and the overload recovery. Increasing one capacitor does not simply add bass: it can prolong a grid-current-induced shift.

Capacitance without double-counting

JJ publishes grid–cathode / grid–anode capacitances [2]: **1.6 / 1.7 pF** (AX7), **2.2 / 1.5 pF** (AT7), **1.9 / 1.63 pF** (AU7). The reduced input capacitance is evaluated from the accepted operating point:

$$C_{\text{input}} = C_{\text{gk}} + C_{\text{ga}} \left[1 + g_{\text{m}} / (g_{\text{p}} + 1/R_{\text{plate}} + 1/R_{\text{next}}) \right]$$

Here g_{m} and g_{p} are the fitted tube-current derivatives and R_{next} represents the next load. This uses **loaded gain**, not an arbitrary fraction of a tube's nominal μ . Each plate also has an **8 pF** lump including output capacitance and inferred wiring. It is an engineering allowance, not a measured chassis value.

The reduction is quasi-static. It omits exact frequency-dependent grid–plate feedback and charge redistribution when the operating point changes. It is more defensible than stacking conservative fixed capacitances, but is not a full physical capacitance matrix.

Tube laws and reference evidence

All voltages below are relative to cathode and currents are in amperes. The curve-validation API and audio path use the same laws. Manufacturer-curve fits are **typical-device models**, not an account of every sample, aging state or production tolerance.

Triodes

The implementation uses a Koren-family softplus law [2, 9]:

$$\begin{aligned} q &= K_p [1/\mu + V_g / \sqrt{K_{vb} + V_p^2}] \\ E &= (V_p/K_p) \log(1 + \exp(q)) \\ I_p &= 2 E^x / K_{g1}, \text{ for } V_p > 0 \end{aligned}$$

The code evaluates softplus safely for very large and small arguments and provides analytic current derivatives for the solver and loaded Miller estimate. Its μ is a **fit coefficient**, not necessarily the datasheet's measured small-signal amplification factor.

Valve	μ	x	K_{g1}	K_p	K_{vb}
12AX7	110.038	1.05000	836.222	653.951	8891.63
12AT7	71.271	1.06979	345.921	285.761	4148.76
12AU7	19.045	1.15058	1223.490	99.167	510.03

These values are fitted to digitized JJ plate-current traces. Capacitance data are separate reference inputs; surrounding resistors are separate circuit choices. Grid-current resistance values and the simplified positive-grid law remain approximations for these triodes.

Pentodes and beam tetrodes

Emission uses a softplus function driven by grid and screen voltage. A smooth collection factor $V_p / \sqrt{V_p^2 + k_{nee}^2}$ supplies the plate knee, with an additional fitted plate-slope factor. The knee also varies with excitation. Screen current has a fitted base emission term plus an extra component near low plate voltage; the 6L6GC and KT88 fits include screen curvature.

Plate curves use JJ EL34, EL84, 6L6GC and KT88 data. Screen references use JJ for EL34/EL84, GE for 6L6GC and GEC/MOV for KT88 [3–5]. Screen supply droop is part of the running circuit, so screen current is not a constant percentage of plate current.

Positive-grid conduction uses $u = \max(V_{gk} - 0.15, 0)$ and $I_g = u(1 + 0.035u)/R_g$. The quadratic form was fitted to published GE 6L6GC positive-grid curves. Its reuse for the other tubes is an **engineering approximation**, not separate screen/grid measurements for every device.

The full-precision tables, digitized points, holdout partitions and DC cases are included in the source archive. Secondary-emission structure, thermal drift of every tube coefficient and exhaustive positive-grid/cutoff behavior are not independently identified. See page 14 for retained accuracy misses.

Tone, inversion and grid memory

Tone and inverter input

The passive Tone pole has $f_c = 650 \times 4000^{\text{Tone}}$ Hz, where Tone is normalized 0–1. The discrete update is $y = a \cdot y_{\text{old}} + (1-a) \cdot x$, with $a = \exp(-2\pi f_c / f_s)$. It is the exact held-input first-order RC update; the broad taper is a **product engineering choice**, not a named vintage tone stack. There is no output shelf or inverse-frequency compensation.

The inverter input capacitance uses **5 pF** inferred wiring plus half the selected tube's grid–cathode capacitance and a gain-dependent Miller term. Its gain estimate uses half of $g_m / (g_p + 1/82 \text{ k}\Omega + 1/220 \text{ k}\Omega)$, consistent with the long-tailed pair's divided gain [8]. It remains a reduced admittance rather than a measured whole-inverter transfer.

Long-tailed pair

Two triodes share a **22 k Ω** tail to **–45 V**, with **82 k Ω / 100 k Ω** plate loads. The unequal loads and shared cathode permit unequal excursions and nonlinear differential behavior. The selected triode law is used for both halves.

The compensation capacitor is represented as **94 pF from each plate to ground**, equivalent to **47 pF between plates** for purely differential motion. The 47 pF reference comes from the Fender 5F6-A circuit [10]. The reduced two-shunt form adds nonphysical common-mode loading when strongly unbalanced; this limitation is retained and documented.

Power-grid coupling and blocking

Each branch has **100 nF** coupling, a **220 k Ω** leak, **5.6 k Ω** stopper and **35 pF** effective grid capacitance. The latter is an engineering lump. The coupling solver includes the positive-grid law, not just a hard voltage clip. When the grid conducts, it draws charge from the capacitor; that changed charge subsequently shifts grid bias and recovers through connected resistances.

Consequently, a bass burst can alter the response to the note that follows it. Blocking is a circuit state, not an envelope-triggered effect. Its duration depends on excitation, bias, resistance and charge; there is no universal fixed release time.

Negative feedback and Presence

Feedback uses an overall factor $0.22 \times \text{Feedback}$. Presence modifies its high-frequency branch around a **1.8 kHz** pole, with depth scaled by $0.85 \times \text{Presence}$. Increasing Presence changes loop response rather than inserting a treble shelf after the amplifier. Its effect can be small if Feedback is near zero.

The feedback predictor uses the local output-stage slope, including the magnetic tangent, to approximate same-sample interaction. The final output-stage solve is nonlinear. This and the screen/supply partition keep the computation bounded; they are not a fully simultaneous SPICE solve of every node.

Output valves and transformer

Matched output configurations

The following are **design targets / circuit choices**, not promises of measured idle B+ or bias. Actual rails change with loading, rectification and control settings. The screen resistors are local to each output valve.

Family	Nominal supply	Nominal bias	Plate-to-plate load	Screen R
EL34	430 V	−35 V	4.2 kΩ	1 kΩ
6L6GC	450 V	−44 V	5.6 kΩ	470 Ω
EL84	320 V	−12.5 V	8.0 kΩ	1 kΩ
KT88	470 V	−59 V	4.2 kΩ	470 Ω

The transformer primary voltage, plate current and reflected load are solved together. Copper loss contributes unequal plate drops. Local screen voltages respond to screen current through resistors and a small **1 nF** numerical/electrical storage approximation. Output plate current and screen current separately load their respective supply rails.

Transformer equivalent and provenance

Hammond's published winding, inductance and response information is used for two representative families [6]. The ideal turns ratio is $\sqrt{(R_{aa}/8\ \Omega)}$. An 8 Ω nominal impedance is not treated as an 8 Ω resistor at every frequency.

Equivalent	1750N-derived family	1750PA-derived EL84
Published reference primary	3.2 kΩ	8.4 kΩ
Magnetizing L	18.3 H × Raa/3200	70.30 H × 8000/8400
Leakage L	13.41 mH × Raa/3200	66.18 mH × 8000/8400
Half-primary DCR	41.74 / 43.14 Ω × Raa/3200	298.94 / 335.45 Ω × 8000/8400
Secondary R	0.4 Ω	0.98995 Ω
Effective primary shunt C	1.59566 nF at 4.2 kΩ; scales inversely with Raa	450 pF

The inductance/copper starting values are reference-backed, but scaling to the selected load is an **inferred equivalent**. The N-family capacitance is **fitted to the published terminal curve**. The PA capacitance remains an **engineering approximation**. No manufacturer winding geometry or exact distributed capacitance is claimed.

Magnetizing inductance is not constant in the nonlinear core. Its nominal value constrains the small-excitation reluctance; strong flux excursion produces increasing magnetizing current and falling incremental inductance. Leakage, winding resistance, core current and the speaker branch remain connected during saturation.

The published response fixture uses resistive loading. Comparing that fixture directly to the complete amplifier's reactive-speaker voltage would confound different loads; those tests are kept distinct.

The magnetic state

The output core uses a scalar effective-field Jiles–Atherton variant informed by published bulk-ferromagnetic formulations [11–13]. The model stores flux linkage λ , field H , total magnetization M and irreversible magnetization M_{irr} . It is continuous between accepted circuit samples; a rejected solver candidate cannot change magnetic history.

$$\begin{aligned} B &= \lambda / (N A) = \mu_0 (H + M) & H_e &= H + \alpha M \\ M_{an} &= M_s [\coth(H_e/a) - a/H_e] \\ M &= c M_{an} + (1 - c) M_{irr} \end{aligned}$$

Irreversible motion is permitted only toward the anhysteretic magnetization in the current field direction. For an accepted effective-field step of magnitude s , the active update is $M_{irr_new} = M_{irr_old} + s/(k+s) \times (M_{an_new} - M_{irr_old})$. If the directional gate is inactive, the previous irreversible state is retained.

What the power tubes must supply

$$\begin{aligned} I_m &= [l_e H + gap \cdot B / \mu_0 + l_e \cdot keddy \cdot dB/dt] / N \\ keddy &= lamination_thickness^2 / (12 \cdot resistivity) \end{aligned}$$

The electrical output equation includes this magnetizing current, primary shunt-capacitor current and reflected load current. Its derivative supplies the incremental reluctance used in the circuit solve. The audible effect therefore arises from an increasing physical current demand, not a separate bass-dependent waveshaper.

Parameter	Value	Evidence status
$\mu_0 M_s$	2.03 T	Manufacturer-informed GO-steel saturation constraint [14].
a / k	62.5782 / 230.9607 A/m	Fitted to the largest published M130-27s loop.
c / α	0.815115 / 9.73765×10^{-5}	Same fit; not independently measured material constants.
Lamination thickness	0.27 mm	Reference material grade.
Resistivity	$4.8 \times 10^{-7} \Omega \cdot m$	Published GO-steel material value [14].
Area / path, N family	10 cm ² / 25 cm	Inferred effective geometry.
Area / path, PA family	6.5 cm ² / 20 cm	Inferred effective geometry.
Turns and joint gap	Derived	Rating/knee linkage and initial-inductance matching, not measured Hammond construction.

Flux linkage at nominal knee is approximately **1.4735 V·s** for EL34/KT88, **1.7014 V·s** for 6L6GC and **1.2202 V·s** for EL84. Iron changes the saturation-strength equivalent. The inferred joint reluctance is not a claim of a physical spacer in a commercial transformer.

Remanence has no artificial decay timer. Electrical losses and connected states govern post-burst recovery; a rate-independent core can retain magnetization at rest. Magnetic viscosity, temperature, mechanical stress, vector anisotropy and excess loss are not independently modeled. **Small minor-loop accuracy is limited** despite the good largest-loop fit (page 14).

The load, supply and recovery

A speaker the output stage can feel

The reactive model uses the Jensen Electric Lightning 70 8 Ω speaker as an electrical reference [7]. Terminal current drives a mechanical mass/compliance/damping equivalent through the force factor; cone velocity returns a back EMF. A lossy inductance branch models the rising high-frequency impedance more realistically than a single ideal voice-coil inductor.

Quantity	Value	Status
DC resistance R_e	6.33 Ω	Manufacturer value.
Moving mass / compliance	29.9 g / 94 $\mu\text{m/N}$	Manufacturer T/S values.
Force factor Bl	12.7 T·m	Manufacturer value.
Resonance / mechanical Q	95 Hz / 14.77	Manufacturer values.
Series inductance	0.38811 mH	Fitted electrical impedance.
Lossy branch L / R	0.95124 mH / 7.28974 Ω	Fitted above 200 Hz; reduced eddy-current equivalent.

These values describe an electrical/mechanical reduction. Acoustic radiation, cabinet resonances and microphone placement are absent. At high frequency the terminal voltage can rise as impedance rises; that is not an acoustic treble boost.

A shared four-node supply

The reservoir feeds screens through **470 Ω** , the inverter through another **8.2 k Ω** , then the preamp through **10 k Ω** . Storage is the adjustable **16–150 μF** reservoir plus **47 / 22 / 22 μF** at successive rails. A **220 k Ω** reservoir bleeder is included. Stereo doubles storage and halves shared resistances so nominal per-channel operation is preserved while channels interact through the rails.

Rectifier equivalent	Fixed drop	Source resistance before Sag scaling
Silicon	1.6 V	35 Ω
GZ34	12 V	105 Ω
5U4G	32 V	230 Ω

Resistance scales by $0.3 + 1.4 \times \text{Sag}$ per mono equivalent. The source is nominal tube supply + 12 V – rectifier drop. Rectification is cycle-averaged with conduction only when source exceeds reservoir; it is not a resolved mains waveform, heater model or fitted rectifier emission curve. These are declared engineering choices.

Bias supply follows a target depending on Bias and B+, with a **0.16 s** nominal lag. A separate power-related thermal memory evolves on approximately **1.8 s**. Coupling charge, rail droop and magnetic remanence have different causes and recovery laws. Their combination gives a heavily driven burst its history-dependent aftermath; there is no single “sag release” control.

Bandwidth and quality

Version 0.1.3 corrects circuit-rate HF damping and the separate 44.1 kHz FIR passband loss. It retains the tube laws, grid charging, supply, feedback, JA core, transformer/load interaction and fixed compensation. The previous 1x HF rolloff was a numerical artifact, not intended voicing.

Choose a quality setting

New instances start at 1x. Saved sessions retain their chosen quality. All modes report **128 host samples** of latency (2.67 ms at 48 kHz). Use 2x or 4x when lower aliasing matters and CPU permits; 4x is also useful for offline rendering.

Quality	Electrical work per host sample	Default stereo CPU ratio
1x	Two complete circuit steps	0.64
2x	Four preamp / two coupled output steps	0.86
4x	Four complete reference steps	1.14

These are processing-time/audio-time ratios at 48 kHz on a Ryzen 9 3950X. Above 1 exceeds realtime. The measured 1x stress presets stay below 0.77; extreme AU7/EL84 at 2x reaches 1.01. Host, workload and sample rate matter. These are not universal deadline guarantees.

Numerical realization

Fast reactive companions use trapezoidal integration with selective lower-rate prewarping. The 1x linear Tone pole uses exact exponential propagation with quartic input reconstruction. Physical LF coupling, magnetic, mechanical and slow recovery times remain intact. Nonlinear equations and solver tolerances are unchanged; 4x retains the corrected reference trajectory.

The FIR has $128 * \text{factor} + 1$ taps and a 0.47-host-rate cutoff. Its measured 20 kHz roundtrip loss stays below 0.15 dB at both 44.1 and 48 kHz. The former 44.1 kHz upper-band loss is corrected.

At default controls and -36 dBFS, measured 44.1/48 kHz response errors versus 4x stay within 0.46 dB/8.51 degrees at 1x and 0.19 dB/4.42 degrees at 2x. Extreme HF overload differs more: the 1x matrix reaches 2.14 dB and 30.11 degrees in separate cases. Quality primarily trades work and alias rejection, but the modes are not numerically identical.

Evidence and its limits

Reference thresholds were not relaxed to make this release pass. The retained DSP suite reports **14 of 16 passing tests**; the two physical-reference misses below are still failures. Current plugin state/audio tests pass; actual hosted VST3 output agrees with the DSP harness to float32 precision. Earlier pluginval results are retained as historical validation.

Comparison	Result / qualification
Digitized tube plate curves	Held-out normalized RMS errors roughly 1–3.4%; all below the retained 10% plate-curve limit.
EL34 / KT88 screen curves	Approximately 7.84% / 11.72% holdout normalized RMS; these curve tests pass their retained 25% gate.
KT88 nominal screen DC point	33.59% error; fails the 30% gate. Curve and individual DC agreement are different requirements.
Steel largest / middle loops	Field normalized RMS approximately 1.84% / 13.24%. Largest loop fitted; middle held out.
Smallest steel loop	31.99% field normalized RMS; fails the 15% gate. Minor-loop losses are also underestimated.
Speaker impedance holdout	0.524 dB RMS; 1.174 dB maximum magnitude error.
Hammond 1750N passive fixture	0.136 dB RMS across 11 sparse readings using frozen-L continuous RLC. Not a validation of JA at the published 27 dBu level.
Hammond 1750PA fixture	0.857 dB RMS across 8 readings; about 1.35 dB too open at 20 kHz.
Jensen input fixture	+0.108 dB at 20 kHz versus published typical –0.05 dB. Wider published –3 dB edges are not both reproduced.

Numerical and musical checks

The tests exercise deterministic plucks, hard attacks, low-E bursts, chords, sustained notes, HF-rich synthetic material and combined 60 Hz/4 kHz excitation. Telemetry records plate/grid voltages, currents, supply rails, bias, flux, magnetizing current, speaker current and coupling charge. Zero solver failures in a probe does not prove exact hardware equivalence.

The material inversion is bounded to **24** iterations, the output solve to **40**, and inverter Newton to **18** with bounded backtracking. Candidate magnetic state is committed only once after acceptance. Fixed-size audio-path storage, preprepared oversampling paths, and allocation/determinism tests address real-time behavior. They do not make an overloaded CPU meet its deadline.

The 9973 Hz overload probe at 48 kHz gives a folded H8 line at 16216 Hz of approximately –21.11 / –32.70 / –119.58 dBc at 1x/2x/4x. This is one resolved line, not total alias energy. Hard HF-rich drive can alias strongly at lower quality. The source includes the execution-fidelity report, explicit poor matches and reproducible measurement scripts.

The model has **not** been validated against an assembled physical SAK-TUBE amplifier. Published typical curves, sparse transformer graphs, reduced capacitances and inferred geometry leave uncertainty. Complete hardware equivalence, thermal aging, magnetic viscosity and exact acoustic speaker behavior are outside the claims of this release.

Reference and provenance index

The source download contains the detailed fit/validation notes and the implementation. Original manufacturer documents remain with their publishers. The links below identify sources; component brands do not endorse SAK-TUBE.

[1] **Input transformer.** Jensen, JT-11P-1 [datasheet](#). Winding DCR, case capacitance and fixture response; Lm/leakage and lumping are inferred in this implementation.

[2] **Triodes.** JJ Electronic: [ECC83S / 12AX7](#), [ECC81 / 12AT7](#), [ECC82 / 12AU7](#). Plate curves and published interelectrode capacitances. Numerical emission coefficients are fitted here.

[3] **Output valves.** JJ Electronic: [EL34 / E34L](#), [EL84](#), [6L6GC](#), [KT88](#). Digitized plate curves, with EL34/EL84 screen data.

[4] **GE 6L6GC.** [Manufacturer data reproduced by The Valve Museum](#). Screen, positive-grid and representative DC checks.

[5] **GEC / MOV KT88.** [Manufacturer data archive](#). Screen characteristics and operating-condition checks.

[6] **Output transformers.** Hammond 1750N and 1750PA. Winding data, specified inductances and published terminal-response graphs. Core geometry is not published by these sheets.

[7] **Speaker.** Jensen [Electric Lightning 70 specification](#). T/S values and impedance magnitude. Lossy-inductance parameters are fitted to selected impedance samples.

[8] **Circuit analysis.** Merlin Blencowe, [Common Gain Stage](#), §1.22 and [AC-coupled long-tailed pair](#). Loaded gain/Miller and inverter reductions. The source notes include the Brimar RC examples used to motivate AT7/AU7 bias choices.

[9] **Tube-model literature.** Norman Koren, [Improved vacuum-tube models for SPICE simulations](#). Model-family background; this release uses independently fitted coefficients and its documented reduced pentode/screen formulation.

[10] **Inverter compensation.** Fender [5F6-A schematic archive](#), 47 pF plate-to-plate capacitor. The complete SAK-TUBE amplifier is not a 5F6-A replica.

[11] **Magnetic variant.** Szewczyk and Cheng (2018), [paper](#) and [Case 3 / Eq. 8](#); author's [JAmode code/data](#). The released data name the material M130-27s; the paper's M127-27s wording is inconsistent. Data/code notices are retained in the source.

[12] **Irreversible-state formulation.** Ugarte Valdivielso et al. (2024), [DOI 10.1109/TPWRD.2024.3398790](#). Magnetic parameter estimation and equation conventions.

[14] **Steel properties.** [Thyssenkrupp powercore typical properties](#) and [EDCOR M-6 steel](#). Saturation and resistivity context; neither identifies the actual Hammond core grade.

[13] **Audio-rate magnetic modeling.** Holters and Zölzer, [DAFx-16 paper](#). Circuit-model context; its JA convention is not claimed identical to the selected effective-field form.

Open source and practical notes

SAK-TUBE is free software under **GNU AGPL version 3**. Copyright © 2026 SAK Audio. It is distributed without warranty. The full license and dependency notices accompany the binary and source downloads. The license applies to the software, not to music you process with it.

The corresponding source archive includes the current plugin code, calibration tables, tests, editable/rendered interface assets, JUCE 8.0.6 and its bundled dependencies, plus build, installer, website and manual-generation scripts. Third-party files retain their own licenses. Original manufacturer PDFs are linked rather than relicensed or bundled.

Build the current plugin

Install Visual Studio 2022 C++ tools, the Windows SDK and CMake 3.22 or newer. Extract the corresponding-source ZIP and run from its top directory:

```
cmake -S SAK-TUBE -B build -G "Visual Studio 17 2022" -A x64 `
  -DSAK_JUCE_DIR=third_party/JUCE
cmake --build build --config Release --parallel 4
ctest --test-dir build -C Release --output-on-failure
```

The VST3 bundle appears under `build/SAK-TUBE/SAKTube_artefacts/Release/VST3/`. The two documented physical-reference test misses on page 14 are expected in this source revision; do not remove their checks or describe a failing full suite as passing. For DSP-only experiments, configure with `-DSAK_BUILD_PLUGIN=OFF`.

Sessions, automation and reset

The 15 current host parameters retain their stable IDs. Drive's stored host coordinate is 0–42, while its visible musical position is 0–10; the electrical taper depends on the selected preamp family. Use the host's normalized automation interface rather than treating that stored value as input dB.

Schema-4 states save parameters, not arbitrary instantaneous capacitor/magnetic histories. Reopening a session restores a deterministic initialized circuit. Playback then evolves its state from the incoming signal. Older states map into the current architecture; removed historical selectors are ignored, and old waveform behavior is not reproduced.

Troubleshooting

No plugin in the host: check for a 64-bit VST3 host, confirm the entire bundle is in a scanned VST3 folder, then rescan. Install the official [Microsoft x64 runtime](#) if a runtime library is missing. A standalone executable is not a VST3 plugin.

Clicks or CPU overload: increase the host buffer, reduce active instances, or use 2×/1×. At high host rates, 4× can make the actual circuit rate and CPU cost very high. Offline rendering is appropriate when the real-time deadline cannot be met.

Unexpected darkness or too much distortion: check Tone, input level and both Drive/Power controls. A cabinet IR adds its own response. Output changes monitoring gain, not the amount of distortion already generated upstream.

Digital clipping: lower Output and check subsequent plugins. SAK-TUBE does not contain an automatic loudness controller or a musical output limiter. Consult the source's `HF_REVISION.md`, `TUBE_REFERENCES.md`, `PASSIVE_REFERENCES.md` and magnetic-reference notes for deeper measurement detail.