

SAK-IRON 0.3

SAK-IRON combines a driven analog input circuit with a nonlinear transformer equivalent. Three controls take it from nearly clean operation to strong compression and harmonic reshaping. The sound depends on incoming level, frequency and the circuit's stored state.

Version 0.3 follows a measurement-led redesign using driven SSL Fusion recordings. It models an interacting input and magnetic path; it is not a recovered hardware schematic or a claim of exact equivalence on every source.

Start here

Insert SAK-IRON on a mono track or stereo bus. Start with **Nickel**, **IRON 25%**, **MIX 100%**. Turn IRON until the amount of compression and harmonics suits the source. Automatic level compensation follows IRON and CORE. Reduce MIX when you want the original transients alongside a more heavily driven core.

The amount of saturation depends on the incoming level. A quiet stem and a loud master will respond differently to the same preset. Bass-rich signals generally reach magnetic saturation earlier than high-frequency signals.

The lower portion of IRON is reserved for clean operation and subtle density; 50% is not intended to represent extreme core saturation. Settings near 70-85% reach heavy drive, and the final portion deliberately exceeds the captured reference. These are starting guides, not fixed distortion thresholds.

Controls

Control	Range	What it does
IRON	0-100%	Progressively drives the input circuit and transformer. Both stages remain in the same signal path throughout the range.
CORE	Nickel / Alloy / Steel	Selects provisional transformer headroom profiles. Nickel is the measured reference profile; Alloy and Steel provide progressively more flux headroom.
MIX	0-100%	Blends latency-aligned dry input with the automatically compensated transformer output. At 0%, the output is exactly delayed dry audio.

IRON is an excitation control, not a percentage of distortion. Its gain compensation cancels predictable drive gain. A curve calibrated offline for each core keeps representative program levels comparable as you turn IRON. The curve depends only on the controls: there is no signal detector or gain riding. Transients still compress and bass still saturates first. At extreme settings, different source levels and spectra can produce substantial residual loudness differences; automatic compensation does not reconstruct a collapsed fundamental or undo compression.

Nickel, Alloy and Steel retain their familiar names and preset positions. In this version they are headroom variations of the selected circuit model. They do not identify three measured alloys, and they no longer use the original material-specific Jiles-Atherton models. Profile refinement remains open.

BYPASS crossfades to latency-aligned, unity-gain dry input. The amber indication means bypass is active; **AUTO LEVEL** indicates compensation. The input and output meters show sample peaks with a short decay. They are not loudness, true-peak or magnetic-flux meters.

Adjusting the interface

Drag a knob vertically or horizontally. Hold **Shift** while dragging for fine adjustment. Click a numerical readout to type a value. Double-click a control to restore its default. The core selector has three detents and accepts arrow keys when focused. Mouse-wheel adjustment is disabled to prevent accidental changes while scrolling a host window.

Resize the window from 800 x 464 to 1500 x 870 logical pixels. Its proportions remain fixed; text and markings are drawn at the current display scale over Blender-rendered hardware assets.

Factory starting points

Preset	Core	IRON	MIX	Intended use
Quiet Weight	Nickel	25%	100%	Restrained starting point for a bus or individual track
Mix Density	Alloy	38%	100%	A little more excitation across a full mix
Drum Weight	Steel	55%	85%	Core compression with some original attack retained
Vocal Body	Nickel	48%	65%	Parallel coloration on a vocal
Bass Current	Alloy	62%	100%	Stronger low-frequency magnetic drive
Synth Pressure	Steel	76%	100%	Firm compression and harmonic change
Core Collapse	Nickel	90%	100%	Deliberately excessive excitation
Foundry	Steel	100%	100%	The end of the excitation range

All presets use automatic compensation and turn bypass off. Their names describe starting intentions, not a promise of the same amount of distortion on every recording. Parameters and the selected program are saved with the host session. Preset values are unchanged from v0.2, but the redesigned processing changes their sound.

For a bus, start near the lower end and compare bypass before deciding whether the change helps. For transient material, try stronger IRON with MIX below 100%. For bass, compare long notes and note attacks: the core responds to both frequency and stored circuit state. For extreme use, try the last fifth of IRON and reduce MIX to retain more original weight and attack. Extreme drive can raise output peaks despite automatic compensation; leave output headroom.

Host behavior

SAK-IRON supports mono-to-mono and stereo-to-stereo processing, float and double audio, automation and saved state. It reports a fixed **128-sample latency** (about 2.67 ms at 48 kHz). Dry/wet blending and bypass retain that delay so the host can compensate consistently. Each channel has its own circuit state; identical left and right input produces identical output. Unequal stereo signals can change their correlation through independent nonlinear processing.

Oversampling adjusts to the host rate: 16x below 88.2 kHz, 8x below 176.4 kHz, 4x below 352.8 kHz, 2x below 705.6 kHz, and 1x at or above 705.6 kHz. At common rates, internal processing is therefore 705.6 or 768 kHz. IRON and MIX smooth over 20 ms. Core changes briefly pass through aligned dry while the circuit settles; bypass uses a 5-ms transition.

Version 0.3 retains the *drive*, *core*, *mix* and *bypass* parameter IDs and the existing state format. Older states restore those controls, but the same settings sound different because the nonlinear system has changed. The removed v0.1 **OUTPUT** trim remains ignored. Keep an earlier build for projects that need its exact sound; existing v0.2 distribution packages remain historical

builds.

What is modeled

The selected circuit has finite input headroom with slight asymmetry, a stateful nonlinear flux response, and output AC coupling. It was selected using the input-only and transformer-enabled SSL recordings separately, then checked on the complete kick, drum loop and mix. It is a compact terminal model, not a recovered SSL schematic.

The current transformer law is reversible. Flux and coupling-capacitor states provide memory; irreversible hysteresis, magnetic viscosity and detailed distributed losses are not claimed. Candidate hysteresis models remain in the research tools, but did not show a sufficient advantage to select them. The model still differs in harmonic balance and some transient behavior.

Reset clears circuit, filter and delay states deterministically. No random component drift, added noise, adaptive gain rider or automatic history erasure runs during playback. Read the forensic report (corresponding source: [SSL_V2_FORENSICS.md](#)), model assumptions (corresponding source: [MODEL.md](#)), and research (corresponding source: [SSL_V2_RESEARCH.md](#)) for measured matches, remaining limits and the reasons for the architecture change.

Installation and validation

The Windows x64 installer includes the VST3, reference manual and optional standalone application. Close your audio host, run Setup, choose the VST3 folder, then rescan if needed. The default is `C:\Program Files\Common Files\VST3`; documentation and the optional application go in `C:\Program Files\SAK Audio\SAK-IRON`. Remove the installation through Windows Settings → Apps. Setup installs Microsoft's signed Visual C++ x64 runtime if required. The SAK installer itself is unsigned.

For manual installation, extract the portable ZIP and copy the entire `SAK-IRON.vst3` folder to a location scanned by your host. Retain its internal folder structure. The standalone application lets you select an audio device and input/output channels. Windows 10 or later, 64-bit, is required; the build and host validation were run on Windows 11.

Installer, portable ZIP, matching source, licenses and checksums are available at [SAK Audio downloads](#). The build instructions (corresponding source: [../README.md](#)) reproduce the plugin and bench. Production measurements (corresponding source: [../validation/ssl-v2/production](#)) and candidate comparisons (corresponding source: [../validation/ssl-v2/candidates](#)) document the measured scope. Earlier distribution packages remain separate historical versions.

Technical reference: the coupled circuit

The complete wet path runs at the oversampled rate:

```
input × excitation g → asymmetric rail circuit
→ magnetic drive m → nonlinear flux circuit → divide by m
→ 1.75 Hz AC coupling → divide by g → fixed compensation
→ antialias filter → latency-aligned MIX / bypass
```

The input stage represents finite source impedance and opposing exponential rail junctions. For each polarity it solves:

```
abs(x) = y + t × exp((y - rail_for_polarity) / t)
output = sign(x) × y
```

The selected equivalent rail is 0.208541, the junction scale is 0.00383750, and the polarity asymmetry is -0.0193546. These are capture-referred units, not separately identified hardware supply voltages or components. A bounded Newton solve enforces current balance. At small signal, conduction is negligible.

The transformer equivalent integrates flux while solving its current loading:

```

v = u - i(z)
F × dz/dt = v
i(z) = a×z + k×sign(z)×abs(z)^p + c×((z+b)^3 - b^3)

```

Parameter	Nickel reference value	Role
F	0.0003102715	Equivalent flux scale
a	0.002216939	Small-signal current slope
k	0.025	Saturation strength
p	13.83134	High-field knee shape
c	0.003391040	Weak nonlinear reluctance
b	-0.1655373	Asymmetric operating-point term

The current law is monotone. The steep term models increasing current as the core approaches saturation; the shifted cubic introduces weak asymmetry. Alloy and Steel increase F by 15% and 30%. Turns, area, physical B/H and individual source/load impedances are not separately identified.

The implicit trapezoidal update is:

```

z_new + dt/(2F) × i(z_new)
= z_old + dt/(2F) × (u_new + u_old - i_old)

```

Trial evaluations leave old state untouched; state commits after the solve. The current law uses a matching linear continuation beyond $|z| = 4$ to avoid numerical overflow. This remote continuation is not a measured material law. The output coupling pole is discretized with a bilinear transform. Flux and coupling state allow history-dependent response even though the selected constitutive law has no irreversible hysteresis.

Drive, compensation and automation

IRON's input-excitation dB anchors are (0%, -48), (20%, -30), (50%, -8), (70%, -2), (80%, 0) and (100%, +14), joined with monotone cubic segments. Magnetic excitation is $m = 0.05 + 0.95 \times (\text{IRON} / 0.8)^4$, with IRON expressed from zero to one. At 80%, both excitation factors equal one. The lower range introduces input density before severe magnetic compression.

Automatic compensation is a fixed 21-node dB curve for each CORE, interpolated from IRON alone. Calibration used 504 synthetic pink/brown-noise measurements, two levels, two seeds and settling before measurement. No reference recording was used to set this gain table. Maximum compensation is 8.951 dB; at Nickel 80% it is +0.797609 dB. A constant scalar preserves within-setting harmonic ratios, compression and crest factor.

During falling IRON, flux and its input/current history contract by the ratio of new to old input-times-magnetic excitation; coupling history contracts by the input-excitation ratio. Ratios never exceed one. Rising IRON retains physical state. This prevents control movement from amplifying stored voltage or injecting flux energy. It adds control-induced damping: repeated up/down automation can dissipate history. Fixed settings skip this rule, and signal envelopes never trigger it. It is a product automation convention, not an identified hardware mechanism or a live gain rider.

Measurements and limits

Reference analysis separates RAW, the driven analog path without transformer, and the same path with transformer. Absolute levels are preserved. The dataset contains 28 sine triplets at 20/50/100/1000 Hz and -30 through -3 dBFS, plus complete kick, drum-loop and full-mix triplets. Independent tone starting phases prevent recovery of an unambiguous absolute fundamental phase shift.

Measurement	Result and scope
Clean-range THD	Maximum 0.00808% at IRON 0%; 0.06680% at 20%, across supplied levels, frequencies and cores
Clean-range fundamental gain	Approximately -0.048 to $+0.003$ dB
Capture-anchor fundamental error	0.1002 dB RMS, 0.2307 dB maximum across 28 points, Nickel 80% before compensation
Capture-anchor compression error	0.09062 dB RMS relative to each low-level response
Full-program RMS error	Kick $+0.024$ dB; drums $+0.131$ dB; mix -0.026 dB, before compensation
Harmonic mismatch	At 20 Hz / -3 dBFS, H2 is -54.52 dBc versus reference -27.51 dBc
Quiet LF mismatch	At 20 Hz / -30 dBFS, THD is 3.37% versus reference 1.13%
Extreme aliasing	7 kHz / -3 dBFS at 48 kHz, Nickel 100%: -64.47 dBc coherent nonharmonic energy at 16x

The model rounds kick/drum peaks more than the reference and retains somewhat more crest on the full mix. Numerical agreement is not a guarantee for every source. DC-bias and burst-recovery tests characterize the model internally; they do not identify the hardware's viscosity, thermal behavior or DC response.

The release passes DSP regression, independent circuit invariants, plugin state/audio tests, an actual VST3 host and pluginval strictness 10 with GUI tests skipped. Trapezoidal integration has a small measured energy error and is not asserted exactly passive. Oversampling reduces aliases without making extreme saturation alias-free. The forensic report retains poor matches and the exact measurement/build evidence chain.

Source and references

Original plugin code, artwork and documentation are AGPL-3.0-only. The source archive includes matching JUCE dependency source, shared DSP, Blender artwork, build and installer scripts, measurements and research. Reference audio and audio-bearing analysis arrays are not redistributed; rerunning the hardware comparison requires a separately supplied dataset. Synthetic DSP tests can be run independently. Third-party components retain their own notices.

Useful primary references include the [SSL Fusion user guide](#), [TI amplifier macromodel guidance](#), and [Holters' coupled magnetic modeling work](#). The corresponding-source research notes explain their relevance and the limits of identifying a circuit from terminal recordings. Product names identify the hardware used as a reference; no affiliation is implied.