SAK-IRON 0.4.0

Reference manual - 8 November 2026

SAK-IRON combines a driven analog input circuit with a nonlinear transformer equivalent. Version 0.4 improves low-frequency harmonic fidelity, corrects CORE-switch bursts and quiet-signal numerical errors, and substantially lowers aliasing and processing cost. CORE, IRON, MIX, presets and the fixed 128-sample latency remain compatible. Existing projects load, but their sound changes.

The circuit is informed by SSL Fusion recordings. It is a compact terminal equivalent, not a recovered schematic or a claim of exact hardware equivalence.

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

Mono and stereo, float and double audio, automation and saved state are supported. Latency is fixed at **128 host samples** (2.67 ms at 48 kHz), including dry/wet mixing and bypass. Each channel owns independent circuit history.

Oversampling is 16x below 88.2 kHz, 8x below 176.4 kHz, 4x below 352.8 kHz and 2x at higher supported rates. First-order antiderivative antialiasing integrates the input rail response between internal samples. A compensated symmetric FIR removes the averaging kernel's linear droop and half-internal-sample delay.

IRON, MIX and CORE position smooth over 20 ms. CORE interpolates headroom and compensation while retaining the wet path and circuit state. It no longer passes through dry audio. Bypass transitions over 5 ms to exact delayed dry. Transport start/stop/seek flags alone do not reset the circuit; host reset clears circuit, filter and delay state. An abrupt input or host reset can still produce a discontinuity. Universal click-free operation is not claimed.

Stable parameter IDs are `drive`, `core`, `mix`, `bypass`; state schema remains 2. The removed v0.1 OUTPUT trim is ignored. Keep an older build for exact recall. Reset and preparation are deterministic. Audio processing has no heap allocation, blocking lock, file access, random noise or live gain detector.

What is modeled

The loaded asymmetric input rail feeds a reversible nonlinear flux circuit, followed by 1.75 Hz output AC coupling and fixed control-dependent compensation. A weak biased cubic and separate smooth strong knee with finite saturated slope improve low- and high-level magnetic behavior without a harmonic table.

Flux and coupling provide memory. Irreversible hysteresis, remanence, physical core dimensions and a unique source/load circuit are not identified by the captures. Nickel is the measured reference; Alloy and Steel increase effective flux headroom by 15% and 30%, without separate hardware validation.

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. Measured before/after report (corresponding source: OVERHAUL_REPORT.md), model equations (corresponding source: MODEL.md) and release validation (corresponding source: RELEASE_040.md) describe the current build. Historical research documents are labelled by version.

Technical reference

```
input -> interpolation -> input excitation -> antialiased asymmetric rail
      -> magnetic excitation -> implicit flux circuit -> magnetic recovery
      -> 1.75 Hz coupling -> input recovery -> fixed insertion/control gain
      -> compensated decimation -> latency-aligned MIX / bypass
```

Input circuit

For each polarity `s`, the finite-source rail law solves

```
a = abs(input)
r = rail * (1 + s*asymmetry)
t = max(1e-5, thermal * (1 + s*thermalAsymmetry))
a = y + t*exp((y-r)/t)
output = s*y
```

Independent junction slopes are constrained by input-only B recordings. They address the original model's incorrect even-harmonic trend near clipping. This is a macromodel, not an identified SSL amplifier schematic. Five bounded Halley iterations solve the positive Wright-omega form `w+log(w)=(a-r)/t`. Far below the rail, the negligible-leakage linear limit is used. The rail is evaluated by first-order antiderivative antialiasing between successive internal inputs. For positive magnitudes `a`, `y`, its primitive is

`.5*y*y+(a-y)*(y-t)`. A derivative-corrected small-interval limit avoids subtractive cancellation. Four scalar history values per channel are reset and rescaled with input excitation. The final FIR compensates the linear kernel's half-internal-sample delay and droop.

Flux and source law

The selected finite-permeability model separates the weak biased cubic from a smooth strong-excitation knee. Let `P(x)=(x+sqrt(x*x+w*w))/2`:

```
F(z) = saturatedSlope * (P(z+b-1)-P(-z-b-1)-P(b-1)+P(-b-1))
i(z) = linearLoss*z + F(z) + cubic*z*(z*z+3*z*bias+3*bias*bias)
D(i) = (i + sourceNonlinearity*sourceThermal*expm1(i/sourceThermal))
       / (1+sourceNonlinearity)
fluxScale * dz/dt = u - D(i(z))
output = u - D(i(z))
```

Here `b=kneeBias`, `w=kneeWidth`. The strong branch approaches a finite slope, instead of an ever-steepening power. The weak cubic and strong branch have separate biases. This improves both weak LF distortion and strong LF even-harmonic behavior. It is an effective terminal law, not recovered B/H. `D(0)=0` and `D'>0`; composition with reluctance admits either a source-drop or an asymmetric material-law interpretation. Remote exponential arguments above 60 continue linearly with matching value/slope. This bounds numerical trials, not a measured extreme-field law. Coefficients are in ModelCalibration.h (corresponding source: ../Source/DSP/ModelCalibration.h).

The selected `finite-joint` experiment fitted the input and transformer jointly, using B input-only evidence and a declared subset of C recordings. A loaded-input imbalance multiplier `q` scales the two input asymmetries. A passive linear flux-rate loss `e` is exactly reduced into `fluxScale *= 1+e`, `insertionGain /= 1+e`. The production header folds both relations once; there is no additional state or runtime correction table. The recordings do not uniquely identify `e` as an eddy-current mechanism. See the report for the training/validation boundary and program fitting.

Implicit trapezoidal integration uses

```
c = dt/(2*fluxScale)
z_new + c*D(i(z_new)) = z_old + c*(u_new + u_old - D_old)
```

The scalar residual is strictly increasing, bracketed between zero and the right-hand side. Newton/bisection commits state once after at most 32 iterations. Termination uses relative floating-point backward error, not an absolute flux floor. Failed convergence is counted. Trapezoidal integration is not claimed exactly energy-preserving or passive under arbitrary discrete excitation.

The analytic law is sampled into 32,769 shared value/slope nodes over [-2,2] during preparation. Monotone cubic Hermite interpolation reduces callback cost; outside that interval, evaluation is analytic. This cache evaluates the equation; it is not fitted to harmonic recordings. An independent analytic bridge and direct residual tests quantify error. Zero-centred coordinates, factored shifted differences and `expm1` avoid quiet-signal cancellation. No noise generator, gate, envelope follower or new signal limiter conceals failed solves.

Flux and coupling provide memory. The selected law remains reversible. Genuine irreversible hysteresis, remanence and rate-dependent loss were not identified convincingly enough to add production states.

Drive, compensation and automation

IRON excitation anchors in dB are `(0,-48),(.2,-30),(.5,-8),(.7,-2),(.8,0),(1,14)`. Magnetic excitation is `.05+.95*(IRON/.8)^4`. Both equal one at 80%, where the uncompensated calibration bridge represents the measured capture conditions.

Compensation uses 21 fixed nodes per CORE, calibrated on 504 synthetic pink/brown references, not hardware recordings. Maximum compensation is 8.9128 dB; Nickel at 80% uses +0.7850 dB. It preserves within-setting compression slopes, crest factor and harmonic ratios; it cannot guarantee equal perceived loudness for every spectrum or restore a collapsed fundamental.

Falling IRON contracts flux and historical derivative terms by the reduction in input-times-magnetic excitation. Coupling and antiderivative input history contract with input excitation. Rising IRON does not expand history. This control-only damping prevents old stored energy being amplified by the falling output divisor. It is not identified hardware physics or a pure coordinate change; repeated modulation can dissipate history. Fixed settings and ordinary audio envelopes never trigger it.

Measurements and limits

All 93 reference WAVs were retained: 28 RAW/input-only/transformer sine triplets at 20, 50, 100 and 1000 Hz, seven levels from -30 to -3 dBFS, plus complete kick, drum-loop and full-mix triplets. Absolute levels are preserved. Independent starting phases prevent identification of absolute inter-capture phase.

The joint fit uses 21 LF input-only tones, 15 LF transformer tones excluding -24/-9 dBFS, one 1 kHz/-6 dBFS loading anchor, and left mix seconds 32-40 with four seconds of warm-up and a 3 kHz analysis-only low-pass. Twelve transformer tones, kick, drums and other mix windows are outside the objective. Repeated inspection makes these development holdouts, not a blind benchmark.

Measurement	0.3	0.4	Reference / scope
20 Hz/-3 dBFS H2	-54.524 dBc	-34.010 dBc	-27.507 dBc hardware; error falls from 27.02 to 6.50 dB
20 Hz/-30 dBFS THD	3.370%	1.061%	1.132% hardware
Fundamental RMS error	0.1002 dB	0.0648 dB	Left, all 28 capture anchors
Compression RMS error	0.0906 dB	0.0357 dB	Left, referenced to each -30 dBFS tone
H2 RMS error	11.082 dB	3.280 dB	Left, reference harmonics above -60 dBc
Extreme-drive alias metric	-64.47 dBc	-99.20 dBc	7 kHz/-3 dBFS, 48 kHz, Nickel, IRON100
CORE-switch excess	14.615 dB	0.031 dB	Peak above both steady endpoints in the CORE sweep

Harmonics use the identical Blackman-Harris, integer-cycle method for both builds; THD sums H2-H40 up to the smaller of 22 kHz and .475 times sample rate. Capture-anchor results exclude additional product compensation. Both channels improve; right fundamental/compression errors are .0645/.0409 dB. The alias metric sums coherent nonharmonic energy over the full output Nyquist band; aliases coinciding with true harmonics cannot be separated by it.

Complete program waveform residuals improve for kick (-22.983 to -25.019 dBr) and mix (-19.110 to -19.550 dBr); drums are nearly unchanged (-18.149 to -18.170). Numerically reconstructed 8x true-crest errors change from -1.358 to -1.467 dB on kick, -1.138 to +.593 dB on drums, and +.894 to +1.268 dB on mix. Thus kick peak compression and mix crest remain weaknesses despite waveform gains. Some intermediate-level even-harmonic phases, H4 balance and 1 kHz/-9 dBFS H5 still disagree. No listening sign-off or exact hardware equivalence is claimed.

At 48 kHz the measured stereo render cost falls about 47-50% across IRON20/80/100, using five warmed two-second renders on Ryzen 9 3950X, Windows x64, Release. These throughput figures are not a portable DAW scheduling guarantee. Two-minute overload/DC/silence testing stayed finite with zero failed solves; the final ten-second peak decayed below 1e-67 FS. Quiet-state backward-error, analytic-cache, local energy convergence and independent rail quadrature tests support numerical stability. The circuit is not exactly discrete-passive and extreme saturation is not alias-free.

Six CTest gates and strictness-10 pluginval cover the DSP and actual VST3 host, float/double, state, latency, mono/stereo, small/irregular blocks and automation. Pluginval GUI tests and manual commercial-DAW listening are not included. The release record ties tested binaries to the distributed installer/source.

Source and references

Original code, artwork and documentation are AGPL-3.0-only. Matching JUCE, shared DSP, build tools, experiments and measurement tables are included in the corresponding-source ZIP. Reference recordings and audio-bearing arrays are excluded; reproducing hardware comparisons needs the original dataset. Third-party components retain their own notices.

The complete overhaul report (corresponding source: OVERHAUL_REPORT.md) includes harmonic and transient plots, rejected alternatives and all remaining limits. The implementation (corresponding source: DSP_IMPLEMENTATION.md) defines state/filter ownership and the validation guide (corresponding source: VALIDATION.md) gives reproduction commands. Antiderivative method: Bilbao et al. (2017). Hardware names identify the measurement reference; no affiliation is implied.