PedalScope Methods

Technical reference: how each measurement is made, and how to check it

PedalScope

Describes PedalScope 1.1 (43)
Published 30 September 2026

Abstract

This document states, for each measurement PedalScope makes, exactly what the shipping code computes: the stimulus, the analysis pipeline step by step with its equations, every parameter that determines the result, and the limits of what the result can claim. Every number in it is generated from a machine-readable methods manifest emitted by the shipped code (the source is open: `https://github.com/PhaseDog/PedalScope`), every method is reimplemented independently in Python from that manifest and the published definition, and a continuous-integration parity test test holds the two implementations to each other and to analytic ground truth. A method change in the app that the Python does not follow is a failing build, not a stale page.

© 2026 Richard Hoge. This document is licensed under CC BY 4.0 (`https://creativecommons.org/licenses/by/4.0/`). The accompanying scripts are licensed under MIT (`https://opensource.org/license/mit`). PedalScope and its mark are not part of either grant.

Contents

1 How to read this document **4**
1.1 The pipeline behind the pages . . . 4
1.2 What this document does not do . . . 4

2 Notation and conventions **4**

3 Harmonic Distortion **5**
3.1 Significance and limitations . . . 5
3.2 The stimulus . . . 5
3.3 The pipeline . . . 6
3.4 Read-time predicates on the result . . . 9
3.5 Parameters . . . 10
3.6 The result . . . 11
3.7 Worked example and parity . . . 12
3.8 Limits and failure modes . . . 14

4 Transfer Curve **15**
4.1 Significance and limitations . . . 15
4.2 The stimulus . . . 15
4.3 The pipeline . . . 16
4.4 Read-time predicates on the result . . . 17
4.5 Parameters . . . 19
4.6 The result . . . 20
4.7 Worked example and parity . . . 23
4.8 Limits and failure modes . . . 24

5 Compression **26**
5.1 Significance and limitations . . . 26
5.2 The stimulus . . . 26
5.3 The pipeline . . . 27
5.4 Read-time predicates on the result . . . 28
5.5 Parameters . . . 30
5.6 The result . . . 31
5.7 Worked example and parity . . . 34
5.8 Limits and failure modes . . . 37

6 Gain Map **39**
6.1 Significance and limitations . . . 39
6.2 The stimulus . . . 39
6.3 The pipeline . . . 39
6.4 Read-time predicates on the result . . . 40
6.5 Parameters . . . 42
6.6 The result . . . 43
6.7 Worked example and parity . . . 43
6.8 Limits and failure modes . . . 49

7 Chord IMD **51**
7.1 Significance and limitations . . . 51
7.2 The stimulus . . . 51
7.3 The pipeline . . . 53
7.4 Read-time predicates on the result . . . 55
7.5 Parameters . . . 59
7.6 The result . . . 61
7.7 Worked example and parity . . . 61
7.8 Limits and failure modes . . . 66

8 Waveform Matrix **69**
8.1 Significance and limitations . . . 69
8.2 The stimulus . . . 69
8.3 The pipeline . . . 71
8.4 Read-time facts on the result . . . 72
8.5 Parameters . . . 73
8.6 The result . . . 73
8.7 Worked example and parity . . . 73
8.8 Limits and failure modes . . . 79

9 Glossary **80**

A Scripts, tests and how to reproduce **91**

Changes since the previous version **93**

1 How to read this document

1.1 The pipeline behind the pages

1. **The manifest.** `analysisdump manifest` (a package tool built on the app's own measurement kernel) emits `Docs/Methods/generated/manifest.json`: one object per measurement, every parameter obtained by calling the shipped function that owns it, never by transcription. Where a parameter is a function rather than a number — a window shape, a gate — the manifest records its name, its defining constants and a one-line rule.

2. **The reimplementation.** `Docs/Methods/py/harmonic_distortion.py` is the method written again, from the manifest and the published definition [1, 3], not from the Swift; each further chapter has its own (`transfer_curve.py` for the Transfer Curve, `gain_map.py` for the Gain Map, which imports the first, `compression_curve.py` for Compression, which imports the first's devices and noise and the transfer module's filters, `chord_imd.py` for Chord IMD, and `waveform_matrix.py` for the Waveform Matrix, which imports the first's devices and noise, the transfer module's filters and series truth, and the compression module's stepped tone). It produces the same table the shipped tool produces.

3. **The parity test.** `test_parity_hd.py` (and, per chapter, its sibling) runs the shipped code and the Python on identical synthetic devices and asserts, at every valid point, that they agree to a stated tolerance and that both agree with analytic truth to a stated tolerance. Tolerances were set from measured distributions and are quoted in §3.7, §4.7, §5.7, §6.7, §7.7 and §8.7.

4. **The figures** are generated by the Python, seeded and byte-stable; the UI captures are rendered headlessly from the app's own chart views by the user-guide pipeline.

5. **The drift check.** CI regenerates the manifest, the fragments and the figures and fails on any difference from the committed copies, so what this document states is what the build computed.

1.2 What this document does not do

It does not describe what any measurement sounds like. A measured difference is a measured difference; whether it can be heard is a question the instrument does not answer, and no sentence here claims otherwise. It names no third-party product in a finding.

2 Notation and conventions

Symbol	Meaning
f_s	the sample rate, Hz
f_1, f_2	the sweep's start and end frequencies, Hz
L	the sweep's rate constant, s: $f(t) = f_1 e^{t/L}$
T	the sweep's actual duration, $T = L \ln(f_2/f_1)$
A	the sweep amplitude, linear, 1 = digital full scale
$\varphi(t)$	the sweep's instantaneous phase, rad
$f(t)$	the sweep's instantaneous frequency, Hz
k	a harmonic order; $k = 1$ is the fundamental
K	the highest order extracted

Symbol	Meaning
f_0	a fundamental frequency, Hz — the excitation whose k-th harmonic is read
$H_k(f)$	the harmonic response of order k on the output-frequency axis (§5)
Δt_k	the time advance of order k's harmonic pulse: $L \ln k$
N	the deconvolution DFT length; M the per-order response DFT length
T_k	the extraction window duration for order k
$h[n]$	the deconvolved impulse response, circular over N samples
dB re input	$20 \log_{10} \lvert H_k \rvert$ — input-normalized: a unit-gain linear device reads $H_1 = 1$, i.e. 0 dB
dBFS	dB relative to digital full scale
dBc	dB relative to a carrier (here, the louder played tone or the fundamental)

Angles are radians. Logarithms are natural unless written $\log_{10}$. Time-domain indices are integers $n = 0 \dots N-1$; "circular" means modulo the buffer length. A raised cosine is $\frac{1}{2}(1 - \cos\theta)$ rising or $\frac{1}{2}(1 + \cos\theta)$ falling.

3 Harmonic Distortion

3.1 Significance and limitations

Harmonic Distortion characterizes how a device converts a single input frequency into output at integer multiples of it: for every fundamental f_0 in the swept band and every order $k \le K$, the amplitude the device emits at kf_0, normalized by the input amplitude. One synchronized swept sine captures the whole band at one drive level; the result is K frequency responses $H_1(f) \dots H_K(f)$ and the derived total harmonic distortion curve.

What it cannot tell you, stated once: (a) anything about behavior at other drive levels — one sweep is one operating point; the Gain Map and Compression measurements own the level axis; (b) anything about intermodulation between two simultaneous notes — Chord IMD owns that; (c) whether a read in an order is content the circuit shaped or noise that landed in that order's analysis window — the harmonic shaping predicate reads that separately (§3.4); (d) anything below the rig's own measurement floor (§3.4); and (e) anything about what the distortion sounds like.

3.2 The stimulus

The stimulus is a synchronized exponential sine sweep [2, 3]:

$$x[n] = A \sin\varphi(t_n), \qquad \varphi(t) = 2\pi f_1 L \left(e^{t/L} - 1\right), \qquad t_n = n/f_s,\ n = 0 \dots \lfloor T f_s + \tfrac{1}{2} \rfloor - 1. \quad (1)$$

Its instantaneous frequency is $f(t) = f_1 e^{t/L}$, so the sweep reaches f_2 at $T = L \ln(f_2/f_1)$. The rate constant is not free: the synchronization condition requires $f_1 L$ to be an integer, so for a requested duration $\hat{T}$

$$L = \frac{\max(1,\ \mathrm{round}(f_1 \hat{T} / \ln(f_2/f_1)))}{f_1}, \quad (2)$$

which is what makes the k-th harmonic generated by a memoryless nonlinearity an exact time-advanced copy of the sweep, advanced by

$$\Delta t_k = L \ln k, \quad (3)$$

and therefore separable in the deconvolved response with meaningful phase. The first $\lfloor t_{\text{in}} f_s \rfloor$ samples are scaled by a rising raised cosine and the last $\lfloor t_{\text{out}} f_s \rfloor$ by a falling one. The fades bound the excited band: below $f(t_{\text{in}})$ and above $f(T - t_{\text{out}})$ the device was driven under a ramp, and reads there are properties of the measurement, not the device.

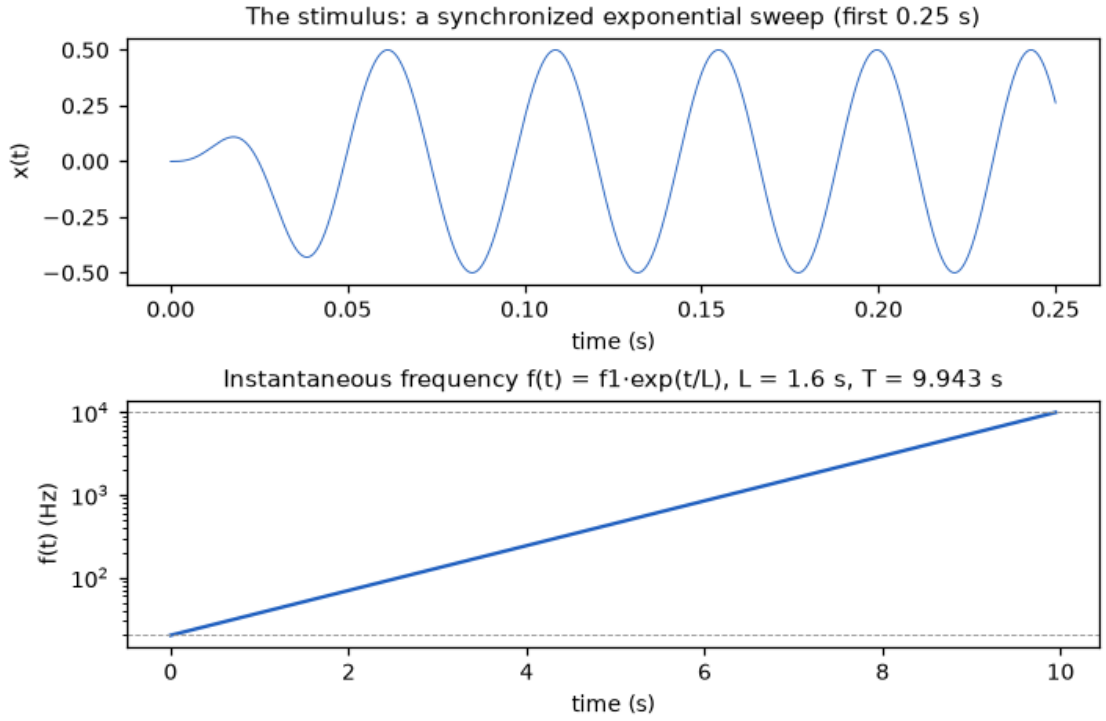


Figure 1: The stimulus, from the reimplementation: the waveform's first quarter second and the instantaneous frequency over the full sweep.

3.3 The pipeline

Given the captured response $y[n]$, time-aligned with the sweep's start (the app measures and removes the loop latency per capture before analysis; the oracle runs at zero latency):

1. **deconvolution with the analytic inverse.** Zero-pad y to N, the smallest power of two not less than the capture length plus $\lfloor 0.1 f_s \rfloor$, and take its DFT $Y[m]$. The inverse filter is the stationary-phase closed form [3]

$$\tilde{X}^{-1}(f) = 2\sqrt{f/L}\,\exp\Big(j\big(\tfrac{\pi}{4} - 2\pi f L\,(1 - \ln(f/f_1))\big)\Big), \tag{4}$$

evaluated at $f = m f_s/N$ for $0 < m < N/2$ and mirrored to the negative bins, tapered to zero below f_1 (a rising raised cosine from $f_1/2$ to f_1) with the DC and Nyquist bins zero. Using the closed form rather than inverting the sweep's own spectrum extrapolates the phase law above f_2, which is what lets harmonics of high fundamentals — whose output frequencies exceed f_2 — deconvolve into clean pulses all the way to Nyquist. The impulse response is

$$h[n] = \frac{1}{N f_s}\,\text{IDFT}\{Y[m]\,\tilde{X}^{-1}(m f_s/N)\}[n], \tag{5}$$

the $1/N$ being the inverse DFT's scaling and the $1/f_s$ converting a sampled signal's DFT to the continuous-time spectrum the analytic inverse expects. A unit-gain linear device yields a pulse of height A at $n = 0$.

2. **Harmonic separation.** Order k's pulse arrives Δt_k earlier than the linear one, i.e. at circular index $N - \mathrm{round}(\Delta t_k f_s)$ (Figure 2).

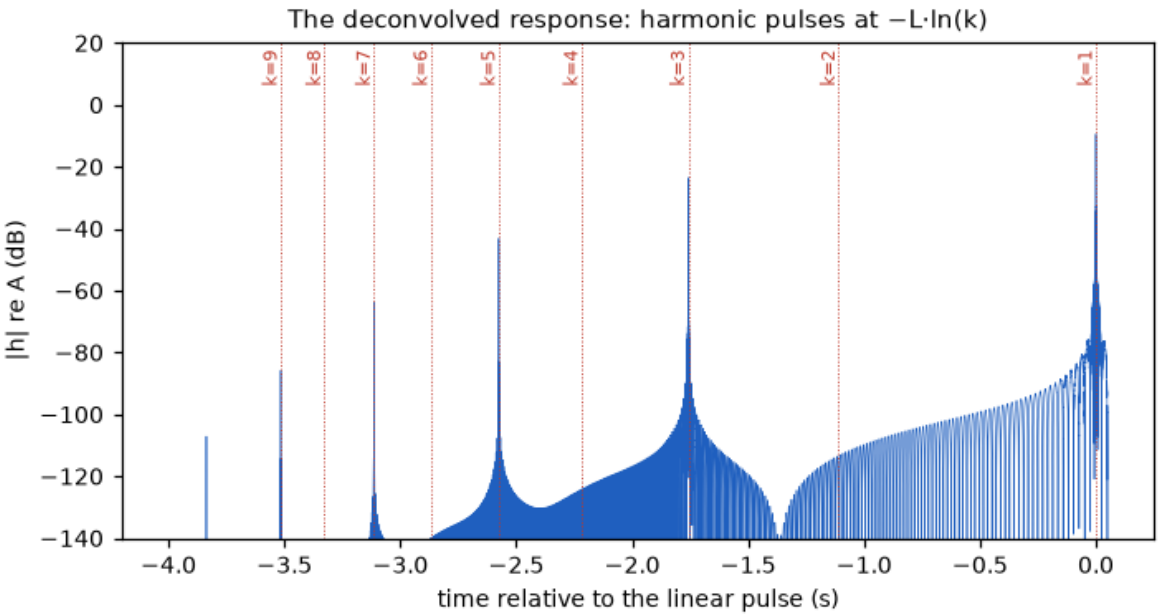


Figure 2: The deconvolved response of the worked example (a symmetric soft clipper), on a negative-delay axis. Odd orders carry content; the even orders' pulses are absent, as they must be for a symmetric device.

3. **Windowing.** Each order is cut out with a raised-cosine window of half-widths ℓ_k (earlier in time, toward order $k+1$) and r_k (toward order $k-1$), in samples:

$$c = \min\left(\tfrac{M}{2} - 1,\ t_{\max} f_s\right), \qquad \ell_k = \left\lfloor \min\left(c,\ \tfrac{(\Delta t_{k+1} - \Delta t_k) f_s}{2}\right) \right\rfloor,$$
$$r_k = \begin{cases} \lfloor c \rfloor & k = 1 \\ \left\lfloor \min\left(c,\ \tfrac{(\Delta t_k - \Delta t_{k-1}) f_s}{2}\right) \right\rfloor & k > 1, \end{cases} \tag{6}$$

with $t_{\max}$ the half-width cap. The window is $w(o) = \frac{1}{2}(1 + \cos(\pi |o|/\ell_k))$ for $-\ell_k \le o < 0$ and $\frac{1}{2}(1 + \cos(\pi o / r_k))$ for $0 \le o < r_k$: unity at the pulse, zero at each edge. An order is emitted only while $\ell_k \ge 16$ and $r_k \ge 16$; the analyzer stops at the first order that fails, so a short sweep simply carries fewer orders. The window duration $T_k = (\ell_k + r_k)/f_s$ is what sets order k's spectral resolution and the correlation width the shaping predicate must clear.

4. **Spectra.** The windowed segment is placed with the pulse at circular index 0 of an M-point buffer (so no linear-phase correction is needed) and transformed:

$$H_k[b] = \frac{1}{A} \mathrm{DFT}_M\{\, h[(p_k + o) \bmod N]\, w(o) \,\}[b], \qquad b = 0 \ldots M/2, \tag{7}$$

with p_k the pulse position. The $1/A$ is the input-normalized convention: $H_1 = 1$ for a unit-gain linear device.

5. **Reading.** H_k lives on the *output* frequency axis. The k-th harmonic of a fundamental f_0 is read at $f = k f_0$, $|H_k|$ interpolated linearly between the two bins adjacent to $f/(f_s/M)$, and reported as $20 \log_{10} |H_k(k f_0)|$.

6. **Validity.** A read is trustworthy only inside the validity bound: $f_1 \le f_0 \le f_2$ and

$$f_0 \le \frac{f_s}{2k+1}, \tag{8}$$

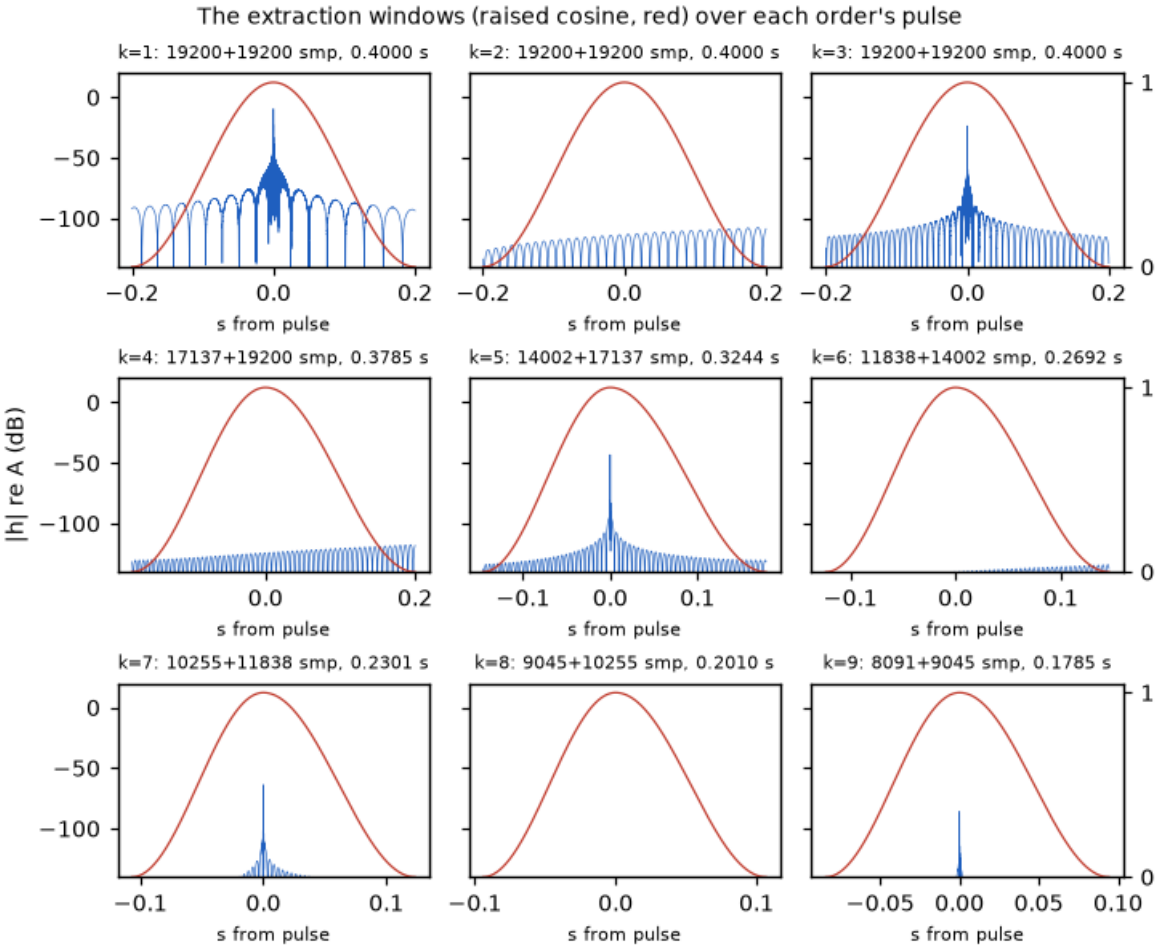


Figure 3: The nine extraction windows over their pulses. Orders 1–3 take the full cap; from order 4 the gap to the neighboring pulse binds.

beyond which the capture folds order $k+1$ onto the read at kf_0 (the alias image bound), and on a loopback-compensated record additionally $f_0 \le f_{\text{cal}}/k$ with f_{cal} the calibration sweep's excited-band top. Invalid reads are excluded from every fit, feature, summary and comparison.

7. **THD.** With the audible band B,

$$\text{THD}(f_0) = \frac{\sqrt{\sum_{k\ge 2,\ kf_0\in B,\ \text{valid}} |H_k(kf_0)|^2}}{|H_1(f_0)|}, \tag{9}$$

THD and not THD+N; a point where the validity bound removed an in-band order is flagged truncated. The curve is sampled on a log-spaced grid of fundamentals from $1.5f_1$ (the first fraction of the sweep is fade-in territory) to f_2.

3.4 Read-time predicates on the result

Three further judgments are computed from the stored spectra when a record is displayed or compared; none is persisted, so fixing a predicate fixes every stored record.

The floor. The rig's own floor comes from the embedded calibration's pooled THD-vs- frequency stamp (one loopback sweep), read at f_0 by linear interpolation on the ($\log f$, dB) plane, then adjusted: per track by $-10\log_{10} N_{\text{tracks}}$ (the pooled figure spread over the harmonic tracks), for drive by the calibration sweep level minus the record's drive, and for sweep averaging by $-10\log_{10}(N/N_{\text{cal}})$ — applied only where the residual the stamp measures is noise (inside the validity bound, above the fade-in edge). A plugin record's operative floor enters as a second, never-averaged bound and the applicable floor is the larger. A read at or below its floor is a bound, not a figure.

Fundamental presence. Every ratio to H_1 (THD above all) is arithmetic unless the fundamental is a measurement: a fundamental sitting at or beyond the analyzer's separation bound below the loudest valid, in-band order is fundamental presence-absent and nothing is divided by it.

Shaping. Anything inside order k's window is reported as H_k, the device's own hiss included. Whether a read is content the circuit deterministically shaped is decided by the coherence of the complex response along frequency: with a lag of $\lceil \gamma/(T_k \cdot f_s/M) \rceil$ bins (clear of the window's own spectral correlation) and n samples,

$$z_i = H(b_i + \text{lag})\,\overline{H(b_i)}, \qquad \Gamma = \frac{|\sum_i z_i|}{\sum_i |z_i|}, \tag{10}$$

which for pure noise is $1/\sqrt{n}$. Γ maps to a signal-to-noise ratio through a pinned Monte-Carlo table, that ratio to the read's own 1-σ dB read scatter (the Rician mean-of-dB uncertainty [4]), and a read is shaped when its ratio clears the bar. The predicate's constants and both tables are in the parameter table below.

Averaging. N back-to-back sweeps in one capture are analyzed separately, each repeat derotated against the first by one fitted delay for all orders (the delay δ maximizing $\text{Re}\sum z\,e^{j\omega\delta}$ over cross-spectral terms sampled per order), and the complex spectra are averaged bin by bin. Magnitudes are never averaged: N reads of a biased quantity average to the same biased number more precisely.

3.5 Parameters

Every value below is read from the manifest, which read it from the shipped code. The per-order table follows.

Parameter	Value
Stimulus	
start frequency f_1	20 Hz
end frequency f_2	10 000 Hz
requested duration $\hat{T}$	10 s
sample rate f_s	96 000 Hz
amplitude A (re full scale)	0.5
rate constant L	1.6 s
synchronization cycles $f_1 L$	32
actual duration $T = L \ln(f_2/f_1)$	9.943372957 s
sample count	954 564
fade-in / fade-out	0.05 s / 0.01 s
excited band (fade edges)	20.635 Hz – 9937.7 Hz
Deconvolution	
inverse-filter taper	10 Hz – 20 Hz
DFT length N for one synth capture	1 048 576
Extraction	
harmonic count K	9
response DFT length M	65 536
window half-width cap	0.2 s
minimum half-width	16 samples (a literal in the analyzer)
bin width f_s/M	1.4648438 Hz
bins per response	32 769
Grids	
distance grid (§11)	96 log-spaced points, 30–10 000 Hz
THD curve grid	128 of 128 requested, 30–10 000 Hz
oracle grid (`synth`)	48 points, 30–10 000 Hz
distance fixed floor clamp	-80 dB
distance minimum resolved share	0.2
audible band	20 Hz – 20 000 Hz
Floor	
calibration analyzed sweeps N_{cal}	1
averaging term at 1 / 2 / 4 / 8 sweeps	0 dB / -3.01 dB / -6.021 dB / -9.031 dB
stamp read (median points)	1
at-floor epsilon	0.1 dB
Averaging	
default / offered sweep counts	4 / 1, 2, 4, 8
pre-roll / gap / tail	0.5 s / 0.5 s / 1 s
delay fit: samples per order	48
delay fit: scan limit / coarse / fine	±8 / 0.02 / 0.0005 samples
Fundamental presence	
analyzer separation bound	76 dB
verdict grid points	128
Shaping	
window samples	32
pure-noise coherence $1/\sqrt{n}$	0.17678
decorrelation factor	2.4

Parameter	Value
shaped SNR bar	6 dB
unshaped scatter	5.57 dB
Summaries	
lexicon version	1.2.2
Oracle (`analysisdump synth`)	
post-sweep pad	48 000 samples
phantom Nyquist gate	0.98 $\times f_s/2$
ground-truth quadrature points	262 144
noise generator increment	`0x9E3779B97F4A7C15`

k	$\Delta t_k = L \ln k$ (s)	window T_k (s)	window (samples)	max f_0 valid (Hz)
1	0	0.4	38 400	9937.69
2	1.10904	0.4	38 400	9937.69
3	1.75778	0.4	38 400	9937.69
4	2.21807	0.37851	36 337	9937.69
5	2.5751	0.324365	31 139	8727.27
6	2.86682	0.269167	25 840	7384.62
7	3.11346	0.230135	22 093	6 400
8	3.32711	0.201042	19 300	5647.06
9	3.51556	0.1785	17 136	5052.63

3.6 The result

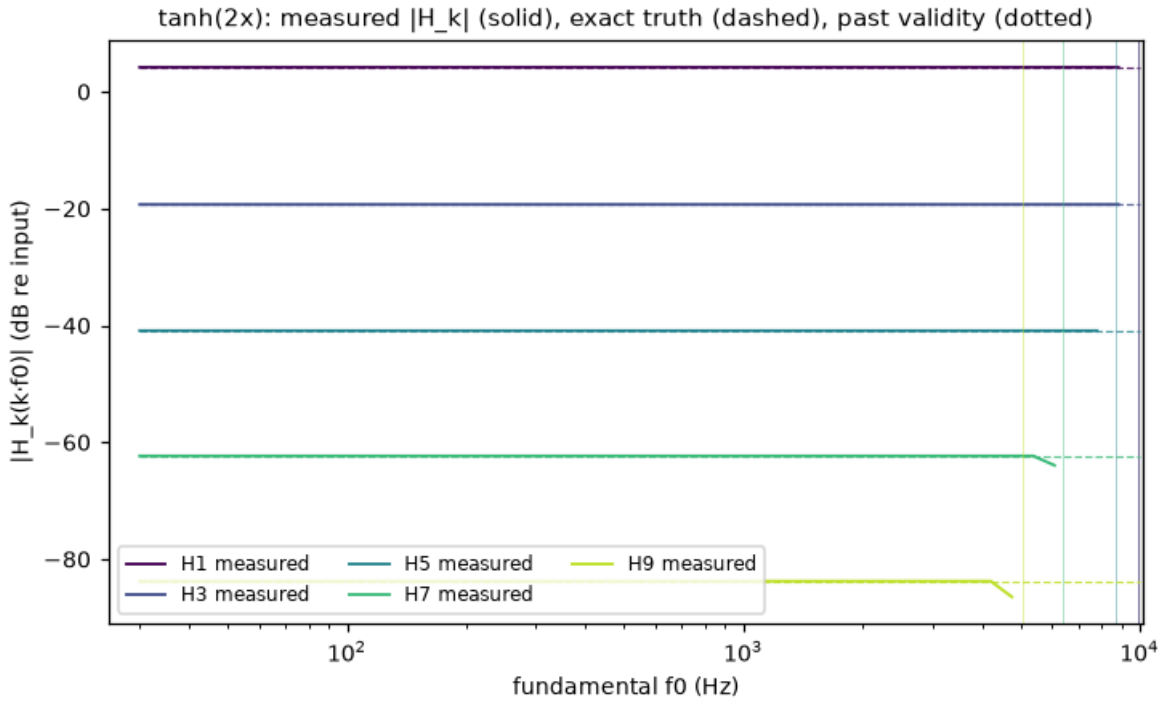


Figure 4: The worked example, $y = \tanh(2x)$ at $A = 0.5$: measured $|H_k|$ for the odd orders against the exact Fourier amplitude of the nonlinearity at the sweep amplitude. Dotted segments lie past each order's validity bound; the shaded strip past the fade-out edge is the excited band's end, where the last grid point lies outside every order's validity bound (it would read the ramp, not the device).

The app's own rendering of the same measurement kind, from the user guide's figure pipeline (a simulated device, rendered headlessly from the chart views, never screenshotted):

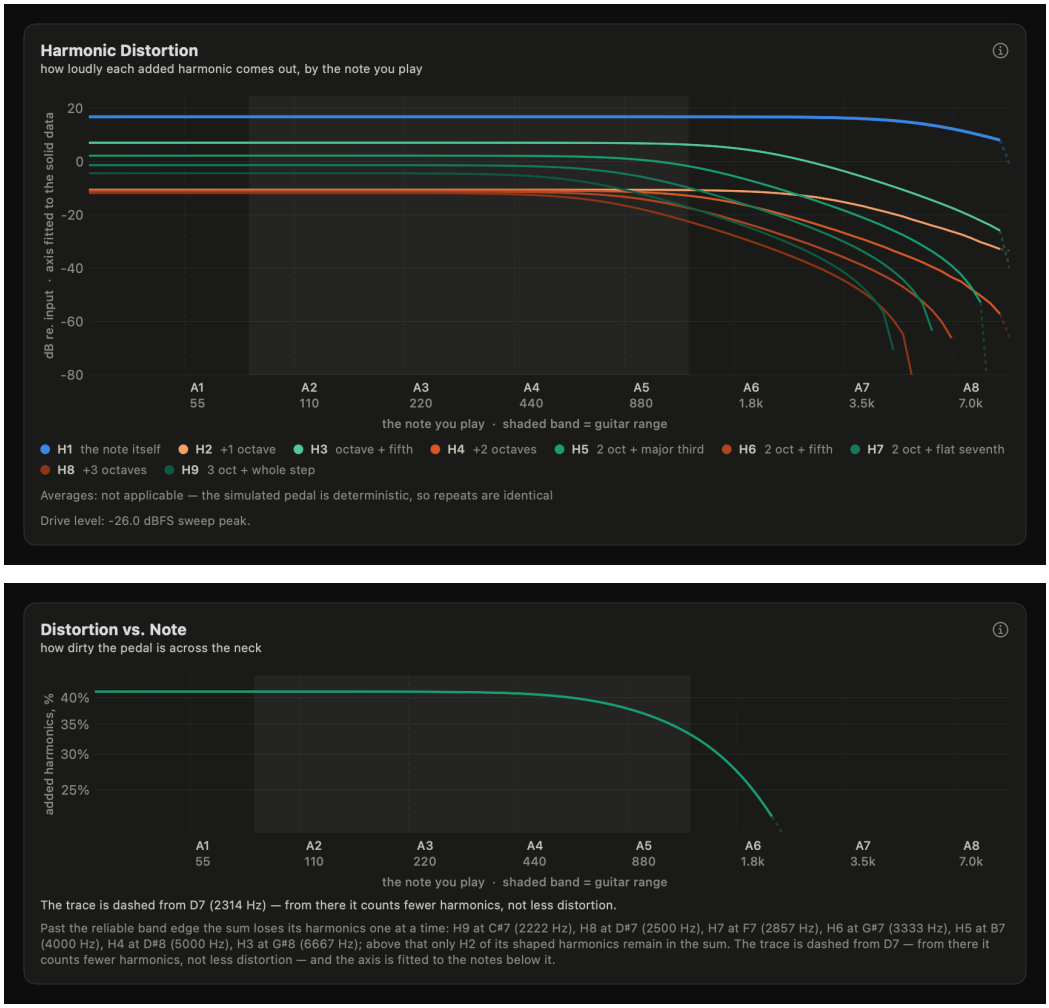


Figure 5: The Harmonic Distortion chart and the Distortion-vs-Note chart as the app draws them, from the guide's simulated fixture.

3.7 Worked example and parity

The parity test runs four synthetic devices through both implementations at one seed with no noise: a phantom capture of three orders (exact by construction), $\tanh(2x)$, a hard clip at ± 0.1 under an amplitude of 0.5, and $x + 0.2x^2 + 0.3x^3$. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. For every valid point inside the excited band it asserts:

- $|\text{Python} - \text{Swift}| \leq 1e-08$ dB on every order that carries content (measured maximum $\leq$ 1e-10 dB on this run; the two implementations differ only by DFT and transcendental rounding);
- both against analytic truth: ≤ 0.01 dB over the whole band for a phantom, and over the lower half of each order's validity band for the oversampled nonlinearities (tanh 0.01 dB, poly 0.01 dB, hardclip 0.02 dB);

- analytically empty orders read at least 50 dB below the case's loudest order on both sides, and the two sides agree on that residue.

The tolerances for the nonlinearities differ from the phantom's because the synthetic device is applied at four times f_s through the same anti-alias chain the app's Simulation Mode uses (the manifest's `synthesis.oversampling`: a Kaiser-windowed sinc whose passband ends at $0.45 f_s$), which is a converter of sorts: it removes the content the device makes above Nyquist instead of letting it fold into the band, and its roll-off lies inside the tops of H_7's and H_9's validity bands, so those half-octaves read the filter rather than the device. Over the full band, $\tanh(2x)$ departs from the bare device's truth by up to 2.71 dB and the hard clip (a near-square wave whose spectrum falls only as $1/k^2$) by up to 2.67 dB, both at the top of H_9. That this is the synthetic device's chain and not the method is shown by each case's *phantom twin*: the same first nine Fourier amplitudes as an alias-free phantom, which the analyzer reads to within 0.01 dB over the whole band.

case	k	expected (dB)	S−P (dB)	S−T band (dB)	P−T band (dB)	S−T full (dB)
phantom	1	0	≤ 1e-10	0.0052	0.0052	0.0052
phantom	2	-10.458	≤ 1e-10	0.000851	0.000851	0.000851
phantom	3	-20	≤ 1e-10	0.000121	0.000121	0.000121
tanh	1	4.2083	≤ 1e-10	0.00492	0.00492	0.00492
tanh	3	-19.292	≤ 1e-10	0.000172	0.000172	0.000172
tanh	5	-40.883	≤ 1e-10	4.2e-05	4.2e-05	4.2e-05
tanh	7	-62.328	≤ 1e-10	4.24e-05	4.24e-05	1.61
tanh	9	-83.757	≤ 1e-10	8.38e-05	8.38e-05	2.71
tanh twin	1	4.2083	≤ 1e-10	0.00515	0.00515	0.00515
tanh twin	3	-19.292	≤ 1e-10	0.000169	0.000169	0.000169
tanh twin	5	-40.883	≤ 1e-10	7.21e-05	7.21e-05	7.21e-05
tanh twin	7	-62.328	≤ 1e-10	3.02e-05	3.02e-05	3.02e-05
tanh twin	9	-83.757	≤ 1e-10	7.73e-05	7.73e-05	7.73e-05
hardclip	1	-11.94	≤ 1e-10	0.00348	0.00348	0.00348
hardclip	3	-21.955	≤ 1e-10	0.0059	0.0059	0.0418
hardclip	5	-27.372	≤ 1e-10	0.0151	0.0151	0.0358
hardclip	7	-31.859	≤ 1e-10	0.0102	0.0102	1.54
hardclip	9	-36.348	≤ 1e-10	0.00482	0.00482	2.67
hardclip twin	1	-11.94	≤ 1e-10	0.00515	0.00515	0.00515
hardclip twin	3	-21.955	≤ 1e-10	0.000179	0.000179	0.000179
hardclip twin	5	-27.372	≤ 1e-10	5.82e-05	5.82e-05	5.82e-05
hardclip twin	7	-31.859	≤ 1e-10	2.04e-05	2.04e-05	2.04e-05
hardclip twin	9	-36.348	≤ 1e-10	6.05e-05	6.05e-05	6.05e-05
poly	1	0.47533	≤ 1e-10	0.00523	0.00523	0.00523
poly	2	-26.021	≤ 1e-10	0.000588	0.000588	0.000588
poly	3	-34.54	≤ 1e-10	0.000136	0.000136	0.000136
poly twin	1	0.47533	≤ 1e-10	0.00516	0.00516	0.00516
poly twin	2	-26.021	≤ 1e-10	0.00106	0.00106	0.00106
poly twin	3	-34.54	≤ 1e-10	9.32e-05	9.32e-05	9.32e-05

For $\tanh(2x)$ the exact amplitudes are $H_1 = 4.2083$ dB, $H_3 = -19.292$ dB and $H_9 = -83.757$ dB re input. A fifth case adds seeded white noise at −80 dBFS and four averaged repeats (the noise stream is reproduced exactly on both sides — same generator, same seeding). On the points the shipped predicate reads as shaped on both sides, the error against truth must sit within 3 times the read's own stated scatter; the last column is that ratio:

k	expected (dB)	shaped points	S−P (dB)	S−T shaped, band (dB)	error / allowance
1	4.2083	42	≤ 1e-10	0.00491	0.022
3	-19.292	42	≤ 1e-10	0.00246	0.036
5	-40.883	41	≤ 1e-10	0.0283	0.41
7	-62.328	38	≤ 1e-10	0.254	0.87
9	-83.757	36	≤ 1e-10	3.9	0.8

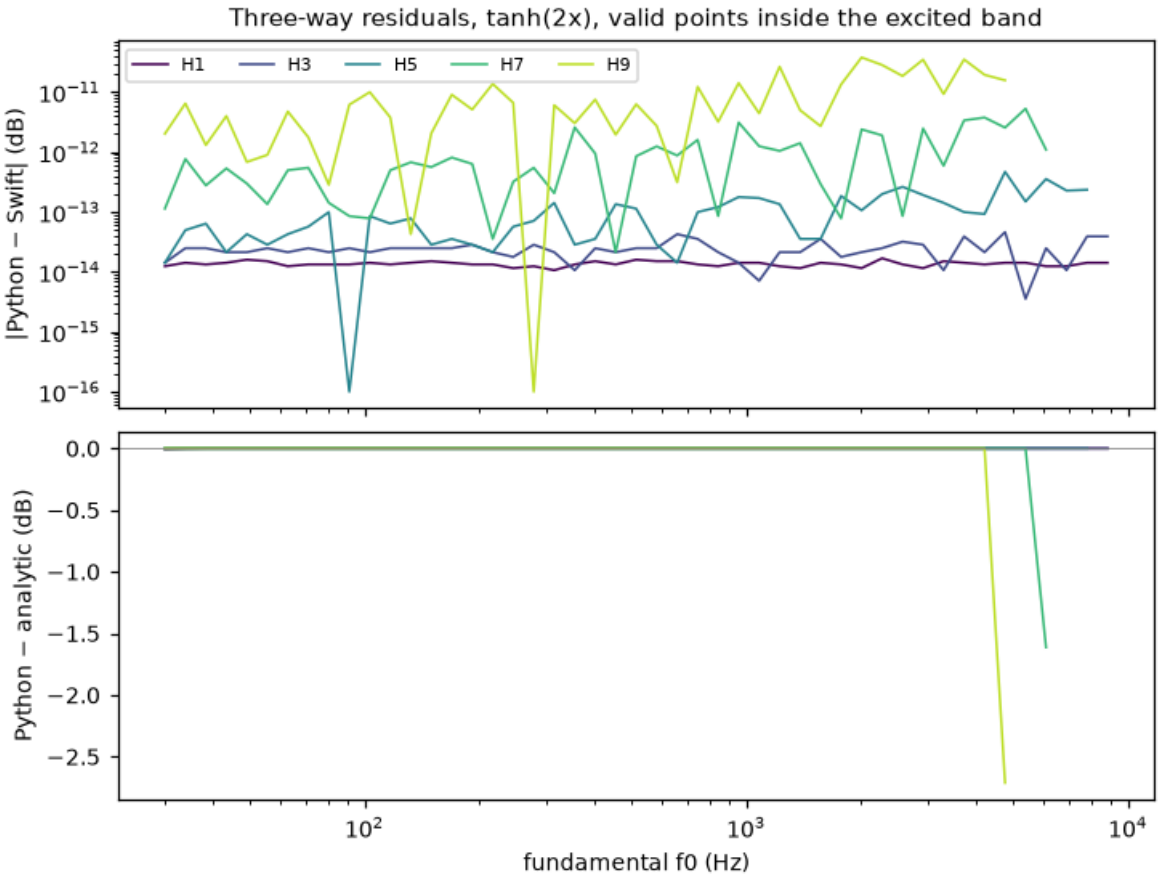


Figure 6: Three-way residuals for the worked example, every valid point inside the excited band. Top: |Python − Swift|, at the level of double-precision rounding. Bottom: Python minus analytic truth, showing the anti-alias chain's roll-off at the top of the highest orders' bands.

3.8 Limits and failure modes

- **The fade edges are measurement, not device.** The last grid point sits on f_2, inside the fade-out, and would read the ramp (the phantom's H_1 there is 11.7 dB low); the excited band's top is therefore part of the validity bound on every analysis — raw, simulated and compensated alike — and that point is outside it for every order. The first fraction of the sweep is treated the same way by every grid's $1.5f_1$ start.
- **Leakage between orders is bounded, and the bound is measured.** An order's read is floored by what the loudest order leaks into its window: worst at the band's low edge, and the value the presence predicate uses is the measured one, posted a hair below what the analyzer demonstrated.
- **A read is not automatically content.** The shaping predicate exists because a symmetric

device's even orders are a read of its noise; every quantity computed from an unshaped read describes the noise.

- **A simulated device's anti-alias chain.** A device applied per sample at f_s folds its content above Nyquist into the band (the hard clip's analytically empty even orders read only 21 dB under H_1 that way); every synthetic device here is therefore applied oversampled through the app's own chain, whose roll-off the parity numbers above quantify for two clippers. Hardware through a converter has the converter's, and the alias-image bound (8) is about the *capture's* folding of the next order, a different mechanism.
- **Short sweeps carry fewer orders.** The 16-sample minimum half-width is a literal in the analyzer, pinned by the tool's guard rather than by the constants census; it binds only for sweeps far shorter than the standard plan.
- **The THD curve's top point.** The shipped curve requests 128 log-spaced points but its top point can land an ulp above f_2 and be dropped by the band guard; the manifest reports the delivered count.

4 Transfer Curve

4.1 Significance and limitations

The transfer curve draws a device's clipping shape directly: while a steady low-frequency sine plays, the instantaneous output is plotted against the instantaneous input over one averaged cycle. A memoryless device retraces a single curve, the static curve, and the figure's bends are its clipping; anything that makes the output depend on where the signal has just been — a shifting bias point, stored charge — opens the figure into a loop. The result is the averaged cycle itself (two vectors of n phase bins), the static curve, and three dimensionless loop areas: the total, the part a linear filter's phase rotation explains, and the part it does not.

What it cannot tell you, stated once: (a) anything about other drive levels — one curve is one operating point, and a clipper whose knee sits above the drive draws a straight line; (b) whether an open loop is memory or filter phase — ordinary tone filters rotate phase and open the loop without any memory, which is why the phase removal exists and why the raw figure's middle band of loop area is hedged; (c) the device's *level* asymmetry once a series coupling capacitor has removed the output's mean, which every pedal's output stage does (§4.8); (d) anything about a rig's own residual alignment, which draws a loop the measurement cannot distinguish from the device's (§4.8); and (e) anything about what the distortion sounds like.

Every measurement described here reads the device at a constant stimulus amplitude after settling. Measurements of the time course of the device's operating point — its recovery after a burst, hysteresis between ascending and descending drive, and its response to a decaying envelope — are in development and will be documented here when they ship.

4.2 The stimulus

The stimulus is a sine at the probe note f_0 (the standard note is E2, at the value the parameter table gives verbatim; other notes are exact twelve-tone equal temperament with A4 = 440 Hz), at a linear peak amplitude A:

$$x[n] = A\,\sin(2\pi f_0\, n/f_s), \qquad n = 0 \ldots \lfloor T f_s \rfloor - 1, \tag{11}$$

scaled over its first and last $\lfloor t_{\text{fade}} f_s \rfloor$ samples by the raised cosine $\frac{1}{2}(1 - \cos(\pi i / N_{\text{fade}}))$, and preceded and followed by silence (the pre-roll and the tail). The app plays this stimulus, measures

the loop latency per capture, and drops the pre-roll before analysis; the oracle renders it through a synthetic device at zero latency. The standard amplitude is the fingerprint sweep's drive (−26 dBFS), so the curve and the other measurement kinds describe one operating point; the worked example below uses the loudest drive the hardware sheet offers, because at the standard drive the reference clippers draw a straight line, which is the point of the sheet's level control.

4.3 The pipeline

Given the tone $x[n]$ and the captured output $y[n]$, aligned by the integer loop latency ℓ:

1. **Phase binning.** With $P = f_s/f_0$ samples per cycle, every sample $i \geq \lfloor s\,P \rfloor$ (s the skipped opening cycles, while the circuit settles) with $i + \ell$ inside the capture is assigned to the bin

$$b(i) = \min\Big(n-1,\ \Big\lfloor n \cdot \big((i/P) \bmod 1\big) \Big\rfloor\Big), \tag{12}$$

   and the averaged cycle is the per-bin mean of $x[i]$ and of $y[i+\ell]$ respectively: $\bar{x}_b$, $\bar{y}_b$, $b = 0 \ldots n-1$. Noise averages; shape does not. Every sample to the end of the tone is binned, the fade-out included — only the start is skipped (§4.8).

2. **The box.** Every area is normalized by the rectangle that bounds the figure (box normalization):

$$B = 4 \max_b |\bar{x}_b| \max_b |\bar{y}_b|. \tag{13}$$

3. **The total loop area.** The closed polyline $(\bar{x}_b, \bar{y}_b)$ is split at every proper self-crossing and the absolute shoelace areas of the resulting lobes are summed — the unsigned lobe area — so that opposite-winding lobes add instead of canceling (a signed integral sees only the fundamental's quadrature and lets a figure-eight report zero). A lobe smaller than a fixed fraction of its own polyline's box is dropped as bin noise. The reported figure is total $= \text{lobes}(\bar{x}, \bar{y})/B$.

4. **The fundamental ellipse.** Both cycles are projected onto $\{1, \cos\theta_b, \sin\theta_b\}$, $\theta_b = 2\pi b/n$:

$$a_x = \tfrac{2}{n}\sum_b \bar{x}_b \cos\theta_b, \quad b_x = \tfrac{2}{n}\sum_b \bar{x}_b \sin\theta_b, \quad a_y, b_y \text{ likewise}, \quad \bar{y} = \tfrac{1}{n}\sum_b \bar{y}_b, \tag{14}$$

   and the fundamental-only trajectory is the least-squares fundamental ellipse, whose area is analytic:

$$\text{linear} = \frac{\pi\,|a_x b_y - a_y b_x|}{B}. \tag{15}$$

   For an input $A\sin\theta$ and an output whose fundamental is $R\sin(\theta+\varphi_1)$ this is $\pi AR\,|\sin\varphi_1|/B$ — a linear filter's rotation of the fundamental, never memory; for a linear device, whose peak is its fundamental's, it reduces to $(\pi/4)\,|\sin\varphi_1|$.

5. **The nonlinear residual.** The ellipse's output is subtracted and the lobes of what remains are summed:

$$\text{nonlinear} = \frac{\text{lobes}(\bar{x},\ \bar{y}_b - \bar{y} - a_y\cos\theta_b - b_y\sin\theta_b)}{B}, \tag{16}$$

   the nonlinear loop area: the harmonic content's own lobes, the headline memory figure. A memoryless clipper's harmonics are in phase with the drive and retrace; harmonics rotated by a filter, or by genuine memory, enclose area here.

6. **The static curve.** The cycle is split at the bins of the input's minimum and maximum into a rising branch and a falling one, each reduced to a strictly ascending polyline in $\bar{x}$, interpolated linearly onto a common grid of inputs from $\min\bar{x}$ to $\max\bar{x}$, and the two interpolants are averaged — per-branch resampling, never binning by input value, whose sample density under a sinusoidal drive is $\propto 1/\sqrt{A^2 - x^2}$ and starves the center.

7. **The memory verdict.** Taken on the nonlinear area (of the compensated figure when the removal is accepted, §4.4): below the first threshold the device reads memoryless-ish, above the second it reads memory, and between them the reading is hedged deliberately — tone filtering *ahead* of the clipper rotates every harmonic by k times the drive's rotation, and a single fundamental cannot separate that from mild memory. Both thresholds are literals in the shipped verdict; the manifest reads them by bisecting it.

Every derived quantity is recomputed from the stored cycle whenever a record is read, so a stored record carries the cycle and nothing else that a fix would have to migrate.

4.4 Read-time predicates on the result

Three judgments are made against the record's paired Harmonic Distortion analysis — the same device, the same setting, the swept measurement of §3 — when one exists; none is persisted.

Orientation. A plugin-path capture is aligned by correlation against the tone and can land half a cycle off, which reverses the figure and, for an asymmetric device, pairs $x(t)$ with $y(t + P/2)$. The orientation resolution reads the cycle's fundamental phase $\psi_1 = \angle Y[1] - \angle X[1]$ (direct DFTs of the two cycles); when it sits within the lock tolerance of $\{0, \pi\}$ the capture is detectably locked, and when the paired H_1's polarity confidently disagrees with $\cos\psi_1 < 0$, the output cycle is advanced by $n/2$ bins and the figure recomputed. Every other case displays the figure as captured and implies nothing. Polarity itself is the energy-weighted mean unit phasor of H_1 over the low band,

$$m = \frac{\sum_b |H_1[b]| \operatorname{Re} H_1[b]}{\sum_b |H_1[b]|^2}, \qquad \text{inverting} \Leftrightarrow m < 0, \quad \text{confidence} = |m|, \tag{17}$$

the real part and never the resultant length: a quadrature response is coherent yet polarity-mute.

Fundamental presence. The memory verdict, the polarity statement and the pairing verdict all key off a fundamental; the fundamental presence reading is taken from the paired record at f_0 (which separates orders far better than one averaged cycle) and from the cycle's own DFT only when the paired reading is undetermined. Where the fundamental is absent the page draws the raw figure and states none of the three.

The phase removal. Harmonics generated *behind* pre-clipper dispersion inherit $k\,\varphi_{\text{pre}}(f_0)$, while an H_1-based rotation would remove $\varphi_{\text{pre}}(kf_0)$; no single shift reconciles the two, so the removal is fitted against the paired record's own branch phases $\angle H_k(kf_0)$, which carry each harmonic's true burden $k\,\varphi_{\text{pre}}(f_0) + \varphi_{\text{post}}(kf_0)$ by construction. The steps:

1. **Fittable terms.** The cycle's harmonic phases are read from its DFT, input-relative: $\text{cyc}_k = \angle Y[k] - k\angle X[1]$. Order k votes when the paired record's validity bound accepts (k, f_0) and, for $k \geq 2$, both $|Y[k]|/|Y[1]|$ and $|H_k(kf_0)|/|H_1(f_0)|$ clear the magnitude floor; its weight is $w_k = \min(1, |Y[k]|/|Y[1]|)$ and its phase difference

$$d_k = \operatorname{wrap}(\text{cyc}_k - \angle H_k(kf_0) - \gamma_k), \qquad \gamma_k = \operatorname{wrap}((k-1)\,\pi/2), \tag{18}$$

   γ_k being the Novák factor between the cycle's convention and the deconvolver's. Fewer than the minimum number of terms means insufficient evidence, never a verdict.

2. **The frame fit.** The two captures each measured their own latency, so a shift between their frames is legitimate and must never be applied as rotation. One shift Δl (samples) is fitted by maximizing the weighted phase coherence

$$C(\Delta l) = \sum_k w_k \cos(d_k - k\,\omega_0\,\Delta l), \qquad \omega_0 = 2\pi f_0/f_s, \tag{19}$$

by a coarse scan over $\pm P/2$ and a fine scan around its maximum (both steps are literals, stated in the parameter table): the frame fit. The residuals $r_k = \mathrm{wrap}(d_k - k\omega_0\Delta l)$ give the pairing residual, the weighted RMS over $k \geq 2$ (the fundamental can always be absorbed into Δl, so consistency lives in the harmonics); above the tolerance the outcome is a pairing mismatch.

3. **The anchored sequential fold.** The frame-mapped branch phase $m_k = \mathrm{cyc}_k - r_k$ is what the paired record predicts for harmonic k in the cycle's own frame. Its burden over the memoryless memoryless lattice $\{0, \pi\}$ is known only modulo π (the multiple removed is the device's own harmonic sign, which the figure keeps), and the folds must come out mutually consistent. The fundamental is anchored: its lattice point is π when the paired polarity reads inverting and 0 when not (the nearest point only when polarity is not confident, and then the sign is quotable only inside the lattice-sign confidence), giving $e_1 = \mathrm{wrap}(m_1 - \sigma_1)$. Each higher voting order, in ascending order, folds against a trend: $k\,e_1$ for the first, and for every later one $a\,k + F(k)$, where $F(k) = \phi_1^u(k) - \phi_1^u(1) + (e_1 - \omega_0\Delta l)$ is the paired H_1's unwrapped phase *shape* anchored at the fundamental's burden (the post-clipper term) and the slope a is refit after each fold over the orders already folded (the pre-clipper term; the fundamental never votes in it):

$$\beta_k = m_k - \pi\,\mathrm{round}\Big(\frac{m_k - \mathrm{trend}_k}{\pi}\Big), \qquad a = \frac{\sum_{j\leq k} w_j\, j\,(\beta_j - F(j))}{\sum_{j\leq k} w_j\, j^2}. \tag{20}$$

The per-order assumption is that the curvature beyond the ramp is under $\pi/2$ — true of the resistor–capacitor networks pedals are built from, and weakest at the lowest order, where it holds across every device class measured.

4. **The all-pass.** The correction is a smooth function of frequency, applied to every bin of the cycle's spectrum at unity magnitude: the deviations $\beta_k - ak - F(k)$ at the voting orders are interpolated by a shape-preserving cubic (Fritsch–Carlson, constant beyond the end nodes), and for $1 \leq k < n/2$

$$Y[k] \leftarrow Y[k]\,e^{-j\theta_k}, \qquad \theta_k = ak + F(k) + \mathrm{dev}(k), \tag{21}$$

with DC and the real Nyquist bin untouched. The compensated cycle is the inverse real DFT and the figure is recomputed from it. Content above the voting orders survives phase-corrected rather than being resynthesized away, which is what removed the truncation ripple an earlier per-order resynthesis drew on hard-clipped plateaus. The correction is *measured* up to the support $k_{\max} f_0$, the highest voting order times the drive note; above it the ramp is extended, F follows H_1's phase as far as its grid holds a value, and the deviation holds its last knot — an extrapolation the page's caption states.

5. **The outcome.** One of five: compensated; a pairing mismatch (no shift reconciles the harmonic phases); insufficient evidence (too few fittable terms); a negligible raw loop (the total under the dead zone, decided after the fit ran so polarity still rides it, and outranking a mismatch — a noise-dominated tiny loop must not wear a warning); and a tripped metric guard, when a removal that passed the residual still grew a *decomposed* area beyond its tolerance (the unsigned total is never compared: the ellipse and the harmonic lobes partially cancel in it). A pairing mismatch has two possible causes the residual alone does not

separate: a session mismatch — another setting or another time, resolved by re-measuring both kinds together — and a structural one, a device whose harmonics have no fundamental-anchored phase relationship (a frequency doubler), which no re-measurement resolves.

6. **Polarity, confirmed.** The figure states polarity only when a fit ran, its residual confirms the pairing, and the fundamental's lattice sign resolved; the robust H_1 statistic supplies the sign, the fit the trust. Absence is honest; a guessed sign is not.

An older correction — rotating harmonic k by $\angle H_1(kf_0)$ alone — remains in the kernel with no consumer, and is not the method this document describes.

4.5 Parameters

Every value below is read from the manifest, which read it from the shipped code; the two verdict thresholds were located by bisecting the shipped verdict, and the scan steps, the bin count, the grid, the polarity band and the analyzer's unused default are literals the manifest carries as rule strings.

Parameter	Value
Stimulus	
standard note / frequency f_0	E2 / 82.4 Hz
standard amplitude A (re full scale)	0.05 (-26.02 dBFS)
tone duration T / pre-roll / tail	1.5 s / 0.2 s / 0.2 s
fade-in and fade-out t_{fade}	0.02 s
skipped cycles (the app's call)	8 (the analyzer's own default is 5, unused)
sample rate f_s (the manifest's)	96 000 Hz
samples per cycle P at f_0	1165.0485
probe notes offered	E2, A2, E3, A3, E4, A4, E5, A5, E6, A6, C7
note-range capture	E2, A3, A4
drive level bounds (hardware / plugin ceiling)	-60 to -6 dBFS / 0 dBFS
Binning and the figure	
phase bins n	256
static-curve grid points	101
lobe floor (fraction of the box)	0.0002
Memory verdict (bisected from the shipped verdict)	
memoryless-ish below	0.05
memory effects from	0.15
Phase removal (§ 6.5)	
dead zone ε (raw total area)	0.01
pairing residual tolerance	0.25 rad
fit magnitude floor (both sides)	0.005 (-46 dB)
minimum fittable terms	3
lattice sign confidence	0.6 rad
metric-guard tolerance (decomposed areas)	0.02
frame-fit scan: coarse / fine step	0.05 / 0.001 samples (literals)
Orientation and polarity	
half-cycle lock tolerance	0.05 rad
polarity confidence threshold	0.35
polarity band	20 Hz – 2 kHz (literals)
Fundamental presence	

Parameter	Value
analyzer separation bound	76 dB
Oracle (`analysisdump transfer-synth`)	
biquad default Q	0.70710678
truth series orders K	200
truth quadrature points	262 144

4.6 The result

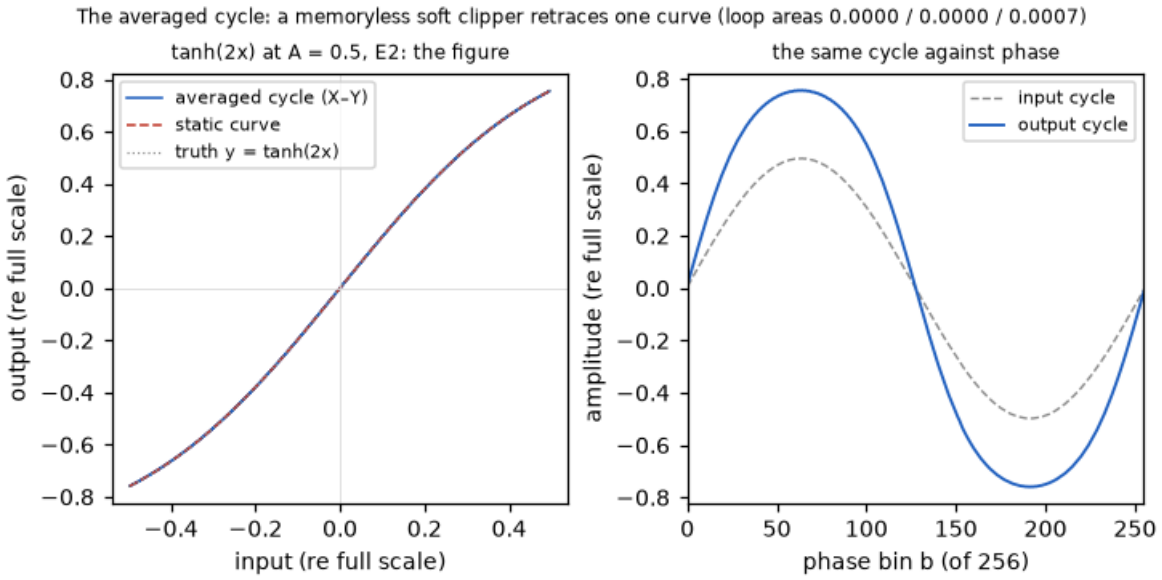


Figure 7: The worked example, $y = \tanh(2x)$ at $A = 0.5$ and E2, no filter: the averaged cycle as the X–Y figure with the static curve and the exact curve overlaid (left; the three coincide), and the same cycle against phase (right). Every loop area is zero to the bin floor.

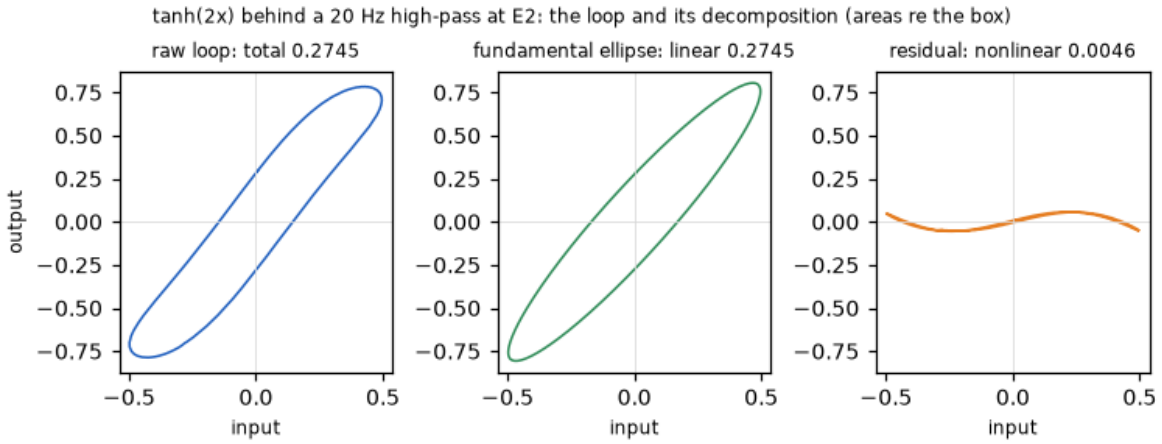


Figure 8: The same clipper behind a 20 Hz second-order high-pass at E2: the raw loop, the fitted fundamental ellipse and the residual after the ellipse's output is subtracted, each with its normalized area. The filter draws the ellipse; the residual holds only the harmonics' own rotation.

The app's own rendering of the same measurement kind, from the user guide's figure pipeline (a simulated device, rendered headlessly from the chart view with the removal on):

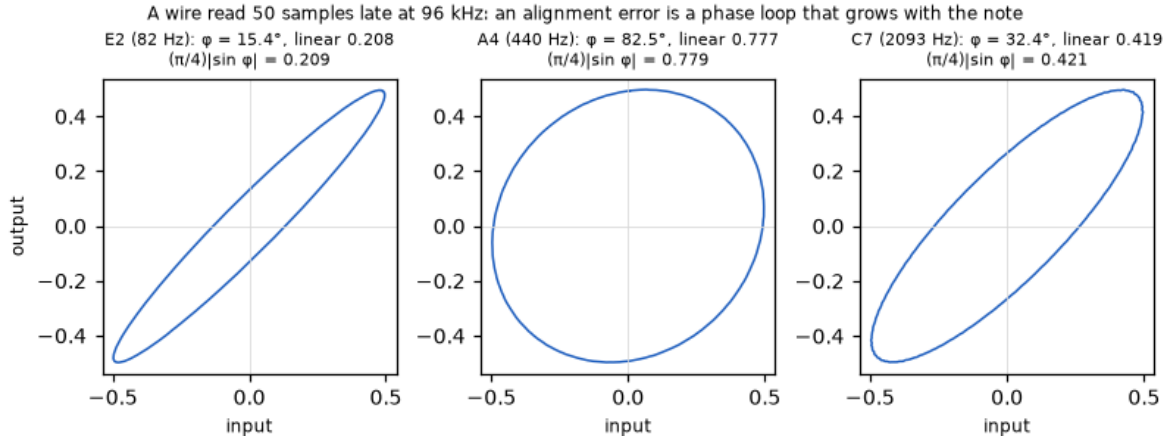


Figure 9: A wire, analyzed 50 samples late at 96 kHz, at three notes: an alignment error is a phase loop that grows with the note. Each title gives the phase the error implies at that note, the measured linear area and its closed form $(\pi/4)|\sin\varphi|$.

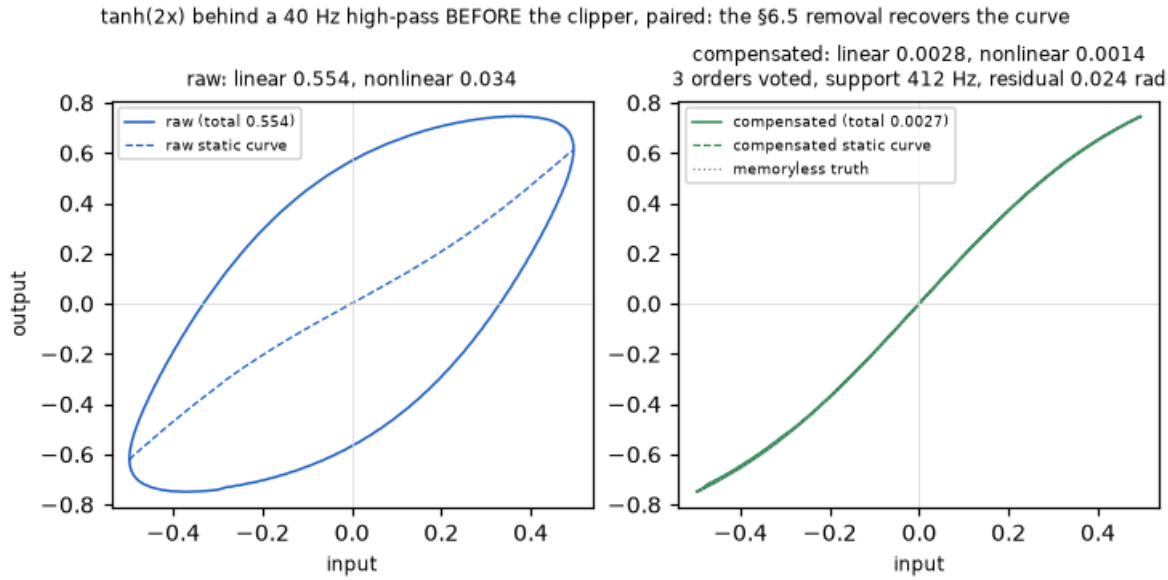


Figure 10: The soft clipper behind a 40 Hz high-pass placed *before* the clipper — the pre-dispersive case an H_1-only rotation cannot handle — paired with its own swept analysis. Left: the raw figure and its static curve. Right: the compensated figure over the memoryless truth, with the orders that voted, the support and the residual.

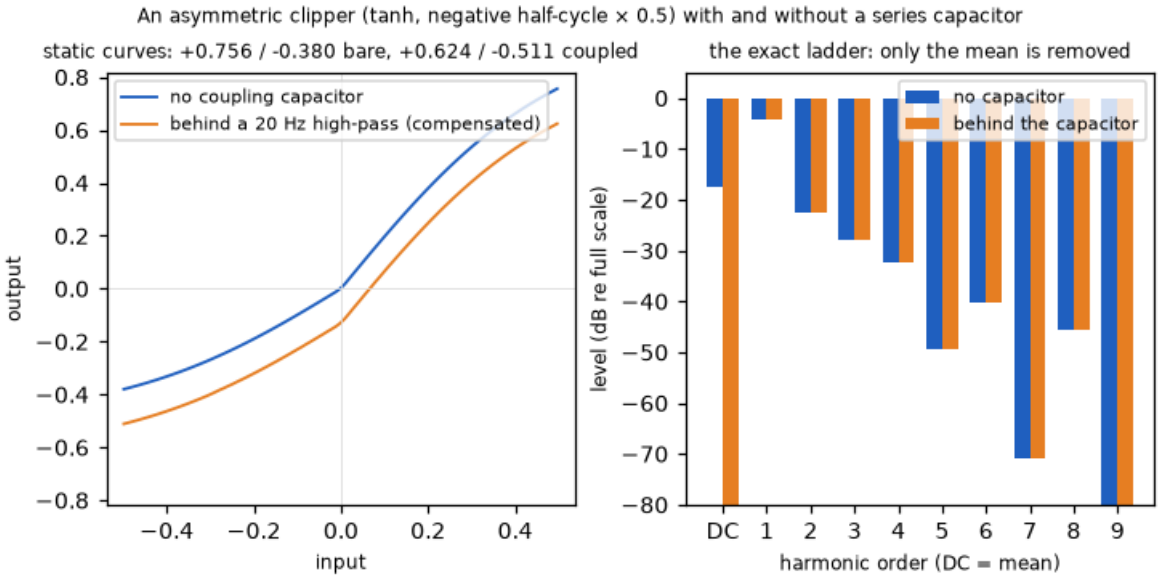


Figure 11: An asymmetric clipper (the negative half-cycle scaled by one half) with and without a series capacitor: the static curve's unequal clip levels are centered by the capacitor while the exact harmonic ladder loses only its mean — the even harmonics survive.



Figure 12: The Transfer Curve view as the app draws it, from the guide's simulated fixture, with the phase removal applied and its caption stating the support and the residual.

4.7 Worked example and parity

The parity test runs 21 synthetic cases through both implementations, each a known device — one of the reference nonlinearities between an optional pre-filter and an optional post-filter, each a second-order (Butterworth) low-pass or high-pass — driven by the plan's own tone and, where paired, measured by the swept method of §3 for its reference. Beside every measured quantity the oracle prints a truth computed by quadrature: for a memoryless device with no filter the curve $y(x)$ itself; otherwise the Fourier series of the clipper's output at the pre-filter's delivered drive, each order scaled and rotated by the post-filter's analytic response and by any planted alignment error, and — for the compensated figure — the same series with every phase returned to the lattice and every magnitude kept. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. The test asserts:

- Swift against Python: every cycle value, static-curve value, compensated value, area, ψ_1, residual and fit term within 1e-09 (measured maximum $\leq$ 1e-10 on this run), the frame shift within 1e-09 samples (measured $\leq$ 1e-10), and every discrete reading — class, flip, polarity, tile, verdict, presence — identical;
- the fundamental ellipse against its closed form $(\pi/4)\,|\sin\varphi_1|$ on a linear device: a ratio within 0.01 of one at the plan's tone (measured deficit at worst 0.73 %) and within 0.001 at ten times its length (measured 0.05 %);
- the raw static curve against the truth, as a fraction of the output peak: 0.01 for smooth memoryless devices (measured 0.003 on the soft clipper, 4.4e-05 at the standard drive), 0.05 for the hard clip (measured 0.027: a 256-bin cycle rounds a corner by about an eighth of a bin's input span), 0.05 for an open ellipse (the branch-mean curve of an open loop is least trustworthy where the branches meet); and a memoryless device's nonlinear area under 0.005 (measured 0.002);
- every paired filtered device compensates, finds the planted frame shift within 0.5 samples (measured 0.39), keeps its residual under 0.05 rad (measured 0.0237), recovers the memoryless truth to 0.01 of peak on the static curve (measured 0.0055) and 0.02 on the cycle, and reads both compensated areas under 0.005 (measured 0.0028).

The removal's parity is established: every paired filtered case compensates on both sides and the two implementations agree on the compensated figure to the Swift-vs-Python tolerance.

The memoryless cases, at the loudest hardware drive and (one case) the standard drive, where the soft clipper is within its linear region:

case	A	total	linear	nonlinear	static vs truth (of peak)
a-tanh	0.5	$\leq$ 1e-10	5.95e-06	0.000719	0.00298
a-tanh-plan	0.05	$\leq$ 1e-10	4.52e-08	0.000612	4.41e-05
a-hardclip	0.5	0.00132	9.68e-05	0.00196	0.0266
a-asym	0.5	$\leq$ 1e-10	0.000119	0.000458	0.00337
a-param	0.5	$\leq$ 1e-10	1.41e-05	0.000939	0.00747

The linear devices — the identity behind the low-pass at three notes and under a 50-sample alignment error at three notes — at the plan's tone and at ten times its length:

case	f_0 (Hz)	φ_1 (°)	measured	$(\pi/4)\lvert\sin\varphi_1\rvert$	ratio
b-lp-E2	82.4	-4.3	0.0586	0.058997	0.99327
b-lp-A6	1 760	-100.2	0.77136	0.7729	0.998
b-lp-C7	2093	-113.4	0.71634	0.72094	0.99362
c-delay-E2	82.4	15.5	0.20771	0.20923	0.99275
c-delay-A4	440	82.5	0.77706	0.77868	0.99792
c-delay-C7	2093	32.4	0.41926	0.42128	0.9952
b-lp-E2-15s	82.4	-4.3	0.058976	0.058997	0.99965
b-lp-A6-15s	1 760	-100.2	0.77283	0.7729	0.9999
b-lp-C7-15s	2093	-113.4	0.72058	0.72094	0.9995

The paired cases: two post-filters and the pre-filter on the soft clipper, the asymmetric clipper behind a DC-blocking high-pass, an inverted device, the half-cycle lock and the unfiltered clipper (whose raw loop is under the dead zone, so the removal is a no-op and the fit still confirms polarity):

case	class	voted	support (Hz)	residual (rad)	Δl (samples)	comp. lin.	comp. nonlin.	comp. vs truth (of peak)
d-tanh-lp	compensated	3	412	0.00322	-0.172	0.00049	0.00054	0.0044
d-tanh-hp	compensated	3	412	0.00945	-0.388	0.00024	0.00066	0.0022
e-tanh-prehp	compensated	3	412	0.0237	-0.272	0.0028	0.0014	0.004
f-asym-hp	compensated	7	659.2	0.0108	-0.102	0.0015	0.0008	0.0055
g-tanh-inv-lp	compensated	3	412	0.00322	-0.172	0.00049	0.00054	0.0044
g-tanh-shift	negligible loop	3	412	0.00833	0.401	–	–	–
g-tanh-paired	negligible loop	3	412	0.00816	-0.0763	–	–	–

Two readings from those tables are worth stating in words. The pre-dispersive case's raw nonlinear area, 0.0343, sits in the hedged band of the verdict with no memory in the device at all, and the removal takes it to 0.0014 — the mid-band hedge in the pipeline is exactly this case. And the half-cycle lock is undone: the captured ψ_1 read 3.144 rad, 0.0026 from the lattice, the paired polarity read non-inverting, and the resolved figure ascends.

4.8 Limits and failure modes

- **The normalization is the box's, and the constant is $\pi/4$ only for a linear device.** The ellipse area is $\pi AR\lvert\sin\varphi_1\rvert/(4A\,y_{\max})$; on a clipped output the peak is not the fundamental's and the figure is $(\pi/4)(R/y_{\max})\lvert\sin\varphi_1\rvert$, which can exceed $(\pi/4)\lvert\sin\varphi_1\rvert$ (a square wave's fundamental is $4/\pi$ of its peak). A reader converting a linear-device figure to an angle uses $\varphi_1 = \arcsin(4\,\text{linear}/\pi)$. On the synthetic single-pole check the shipped code reads the closed form to 0.73 % at the plan's tone and 0.05 % at ten times its length, so a deficit of several percent on a hardware capture is in the measurement chain, not in the metric's constant.

- **The fade-out is binned.** The skipped opening cycles protect the start; every sample to the end of the tone is binned, the last t_{fade} under its ramp, a phase-dependent weighting whose share of the average falls with the tone's length. It is the whole of the deficit above (the E2 ratio moves from 0.9933 to 0.99965 between the two lengths), and it lowers the captured input peak by a few tenths of a percent at the plan's length.

- **A residual alignment error is a loop.** An residual alignment error of e samples rotates the fundamental by $\varphi = 2\pi f_0 e/f_s$ and draws a linear area of $(\pi/4)\lvert\sin\varphi\rvert$ on a wire — growing with the note (Figure 9). The integer latency the app removes leaves at least the fractional part, and a rig can leave more; on a device, that loop is indistinguishable from the device's own linear-phase loop. Two notes measured on one rig let a reader test the

hypothesis: a single e must reproduce both deficits. The nonlinear area is untouched by it on a linear device.

- **The coupling capacitor removes level asymmetry, not harmonic asymmetry.** A clipper clamped at $+a$ and $-b$ at roughly equal duty carries a mean near $(a-b)/2$; a series capacitor at the output strips it, leaving levels of $\pm(a+b)/2$ — equal magnitudes from an unequal circuit (Figure 11). The shape asymmetry remains, and the even harmonics in the paired swept analysis are where the asymmetry is read. Because every pedal has one, a transfer curve essentially never shows the raw asymmetry of an asymmetric clipper.

- **The support ends at the highest voting order.** A harmonic too quiet to clear the magnitude floor on either side does not vote (the soft clipper at this drive seats 3 orders, the asymmetric clipper 7); above the support the correction is the extended ramp, the H_1 shape and a held deviation, and a hard-clipped device has content far above it.

- **The branch phase is read at the nearest bin.** The paired record's phase at kf_0 is the nearest bin's, so a read can be off by up to half a bin width times the filter's phase slope, which the frame fit absorbs as a fictitious shift of a fraction of a sample (measured up to 0.39 samples on the parity cases) and leaves in the higher orders' residuals.

- **The bins round a corner.** With n bins per cycle a hard clip's corner is averaged over a bin's input span, and the static curve departs from the exact curve there by a few percent of the clip level; the smooth devices read under one percent of peak.

- **The static curve of an open loop.** The branch-mean estimate is least trustworthy where the two branches meet, at the input's extremes; on a wide ellipse it departs from the true line by a few percent of peak there. The app draws the static curve dashed on every figure; on an open loop its ends are where the estimate is least trustworthy.

- **A declined removal names its class, not its cause.** The residual separates a genuine pairing from a mismatch; it does not separate a session mismatch from a structural one, and the two need different remedies (§4.4).

5 Compression

5.1 Significance and limitations

A compression curve — the app's "touch response" — is one tone at the probe note stepped through a probe grid of input levels from quiet to loud, and at each step the output level of the fundamental and the total harmonic distortion at that note, read on one snapped window by a one-bin estimator. It answers what the sweep of §3 cannot: where, as the drive rises, one decibel in stops being one decibel out — the knee of a hinge fit through the whole curve, a point with a resolution band or a bound at the floor — and where, going down, the distortion falls under a threshold (the cleanup reading). This instrument is NOT the sweep. There is no deconvolution and no extraction window; the estimator is its own (§5.3), and the chapter owns its equation.

What it cannot tell you, stated once: (a) anything at a note other than the probe note — a circuit whose tone stack sits ahead of its clipper compresses differently across the neck, and one curve is one pitch; (b) anything between two steps — the knee is interpolated by a fit whose breakpoint is free to land between rungs, but the fit is a summary of the measured points and a feature narrower than a rung is ambiguous (the fine window exists for that); (c) anything about a knee the grid did not reach — a device already compressing at the quietest step has its knee somewhere below it, and the answer is a bound, never a number; (d) what a THD read near its floor is — the noise bound says where the read is the rig, and the floor class says so per point; (e) the second crossing — the cleanup level is the FIRST crossing from the quiet end, and a device whose THD rises, falls and rises again with drive (§6's level-trend law: THD is not monotone in drive for a soft-knee overdrive) is quoted at its first crossing only; and (f) anything about what the distortion sounds like.

5.2 The stimulus

The grid is N levels evenly spaced in dB from a floor to a ceiling inclusive,

$$\ell_i = \ell_{\min} + (\ell_{\max} - \ell_{\min})\,\frac{i}{N-1}, \qquad A_i = 10^{\ell_i/20}, \tag{22}$$

with an optional fine window merged in: a span $[l, h]$ clamped inside the grid, filled at the fine spacing with its count capped (so a wide window widens its own spacing rather than growing the run), each fine level appended unless a level already sits within a tenth of the fine step of it, and the whole list sorted — so the fine points interleave the lattice. Two drive windows exist, the hardware one anchored to gain staging (the ceiling is the loudest drive that cannot clip the converter; the floor keeps a stated clearance over a staged rig's converter floor and reaches below the old floor so that high-gain knees resolve to points) and the plugin one reaching digital full scale (the parameter table gives both). The floor, the ceiling, the step count and the window are the user's per run, so a record's grid is read from its own points and never assumed.

The stepped tone is one sine at the probe note f_0 whose amplitude walks the grid quiet to loud, the phase continuous across steps:

$$x[n] = A[n]\,\sin(2\pi f_0 n/f_s), \qquad A[n] = \begin{cases} A_{i-1} + (A_i - A_{i-1})\,\frac{1}{2}(1-\cos(\pi\, m/R)) & m < R \\ A_i & \text{otherwise,} \end{cases} \tag{23}$$

n the global sample index, m the index within step i, R the ramp length in samples ($A_{-1} = 0$). Each step is a settle region and measurement window: a settle region of S samples, whose first R carry the raised-cosine ramp from the previous amplitude and whose remainder lets the device settle unread, then a measurement window of M samples. M is the cycle snap: the requested measurement duration rounded to a whole number of periods of f_0 (at least four), so the window

holds an integer number of cycles of the fundamental and of every harmonic order, and the estimator sees no leakage from the tone itself. The last R samples of the whole tone fade to zero over a raised cosine — and that fade lies INSIDE the last step's measurement window, a fact with a consequence §5.8 states.

The assembly, and a mirror. The measurement of record plays a pre-roll of silence, the tone and a tail of silence; the pre-roll's rms is the noise stamp, a per-capture measurement of the loop PLUS the device's own hiss; the analyzer is handed the pre-roll plus the embedded calibration's round-trip latency as its cut. That assembly is not a MeasureKit type: it lives in the app's measurement service (the record's path), by copy in the command-line tool, differently in Simulation Mode (the tone alone — no pre-roll, no stamp, latency zero, its quality's shorter windows and fewer steps; the parameter table gives both qualities) and again in the plugin probe (a development tool). The oracle behind this chapter reproduces the app's assembly once and states it as a mirror, because there is no shipped function to call; a `CompressionPlan` in the kernel is the filed follow-up (§5.8). Nothing capture-time moved for this chapter.

5.3 The pipeline

Given the captured output $y[n]$, the latency ℓ (the pre-roll plus the calibration's measured round trip; on the plugin path the pre-roll alone, since that path measures and removes its latency per capture) and the stamp σ:

1. **The window.** Step i's window is $y[iS' + S + \ell \;\ldots\; iS' + S + \ell + M)$, $S' = S + M$ the samples per step; a window that runs before the capture's start or past its end is skipped, silently.

2. **The gate.** The window's rms $r_i = \sqrt{\sum y^2/M}$; with a stamp, a step with $r_i \le \sigma\, 10^{g/20}$ — within the stated margin g of the capture's own hiss — is the per-step noise gate's: excluded from the curve and counted. Its "output level" would be the hiss, not the tone, and would flatten the quiet end into a phantom knee; the quietest surviving step is then the effective floor, and a bound reads from it.

3. **The estimator.** The amplitude of one frequency f in the window, by a Hann-windowed one-bin DFT with the window's coherent gain compensated:

$$a(f) = \frac{2\left|\sum_{m=0}^{M-1} y[m]\, w[m]\, e^{-j\,2\pi f m/f_s}\right|}{\sum_m w[m]}, \qquad w[m] = \tfrac{1}{2}\bigl(1 - \cos(2\pi m/(M-1))\bigr), \tag{24}$$

   m the window-local index (only the magnitude is read). On the snapped window a pure tone reads its amplitude to rounding; what the Hann buys is robustness to what the snap does not cover — the ramp's tail and a latency cut that is slightly wrong — and §5.7 measures both. A window of eight samples or fewer reads zero.

4. **The point.** $a_1 = a(f_0)$ is the output amplitude. For $k = 2 \ldots K$ in order: if kf_0 reaches 0.95 of Nyquist or exceeds the band's top the loop ENDS (every higher order is excluded); if it sits under the band's bottom the order is skipped; else $a_k = a(kf_0)$ joins the sum. Then

$$\mathrm{thd}_i = \frac{\sqrt{\sum_k a_k^2}}{a_1}, \tag{25}$$

   a ratio of AMPLITUDES over the declared audible band (THD, not THD+N); zero when the fundamental reads zero. At the default note every rate the app offers sums $k = 2 \ldots K$ entirely; the band's top cuts in above $f_0 = 20\,000/K$ Hz. The point carries A_i, a_1, r_i and thd_i; its levels are $20\log_{10}$ of them (with a small literal guard against a zero fundamental) and its gain is their difference.

5. **The curve.** The points sort ascending by input amplitude (a fine window's points interleave the lattice), and the stored curve is the note, the rate, the points, and — as additive fields — the stamp and the gate's count. Nothing derived is stored: the knee, the cleanup and the classes are computed when the record is read, so a fixed estimator retroactively fixes every stored record.

5.4 Read-time predicates on the result

Everything below is computed from the frozen curve when a record is read.

The incremental slopes. Between consecutive points, the incremental slope $s_i = \Delta\text{out}_i/\Delta\text{in}_i$ in dB per dB (1 when the input step is not positive): 1 is clean, 0 a hard limit.

The hinge fit, in four verdicts. With x_i and y_i the input and output levels in dB over $n \geq 4$ points, the hinge fit is

$$y = c_0 + c_1 x + c_2 \max(0, x - k), \tag{26}$$

two lines meeting continuously at a breakpoint k; $s_1 = c_1$ is the pre-knee slope and $s_2 = c_1 + c_2$ the post-knee slope. The breakpoint is swept over a fine set of candidates between the second and the third-from-last probe (both segments keep at least two points), each candidate solved by the normal equations of the columns $[1, x, \max(0, x - k)]$ with a small ridge on the diagonal, and the candidate with the smallest error wins; the whole (k, SSE) profile is kept. Against it stands the single-line fit with its own error SSE_{lin} and slope s_{lin}. Two constants govern: the post-knee threshold τ_s (a knee compresses when $s_2 < \tau_s$) and the pre-knee unity tolerance δ; the probe step $p = x_1 - x_0$ is the LOCAL first rung — on every uniform stored grid it equals the average step exactly, and on a mixed grid with a fine window mid-ladder the average would understate the spacing at the floor. The verdict, in order:

1. *No break.* If the best hinge does not beat the line by the stated factor, the curve is ONE line: clean throughout when $s_{\text{lin}} \geq 1 - \delta$ (no knee), else sub-unity throughout — a BOUND at x_0 with resolution p and both slopes s_{lin} (mild compression everywhere, the knee at or below the floor).

2. *Never compresses.* A break exists but $s_2 \geq \tau_s$: no knee.

3. *The unity guard.* $s_1 < 1 - \delta$: a BOUND at x_0 with resolution p and the fitted slopes. The device is already compressing at the quietest probe, no clean segment anchors the fit, and the fitted corner is an artifact of a curve bent everywhere; the guard is against UNITY, not against the compression threshold — compared against τ_s, every slope between the two passed as a confident point, which is how a drive-max soft-knee overdrive record in the library once reported a false point. The bound carries its fitted pre-knee slope as the evidence the page quotes.

4. *The floor-bound zone.* $k \leq x_0 + p$: a BOUND at k with resolution p — a breakpoint within one probe step of the floor cannot be anchored from below.

Otherwise a POINT at k, whose resolution band is the interval of candidates whose error stays within five percent of the best; half its width is the resolution, never finer than the candidate spacing, and it sets the decimals the level is displayed with. A bound's label reads "$\leq$".

The cleanup reading. The first point already at or above the threshold τ is *never clean*; else the first i with $\mathrm{thd}_i \geq \tau$ gives *cleans up* at the interpolated crossing (the same linear interpolation and literal guard as the Gain Map's (34)); else the raw function returns the top of the probed range, which the cleanup reading reclassifies as *always clean* — the edge of the probe is never dressed up as a property of the device. Only a genuine crossing quotes a level or draws a marker.

The floor, from three sources. The first source is the noise bound: the stamp expressed as a worst-case THD at each point,

$$F_{\sigma,i} = \frac{\sqrt{2}\,\sigma}{a_{1,i}}, \tag{27}$$

because if every bit of the broadband noise power landed in the harmonic bins the THD numerator would read exactly $\sqrt{2}\,\sigma$ — a strict upper bound on what noise alone can contribute, conservative in the refuse-to-call-noise-distortion direction, and one that climbs toward the quiet end as the fundamental shrinks: the THD+N bathtub the chart must name. No calibration stamp enters here; the stepped tone's silent pre-roll is a per-capture measurement of the loop plus the device's own hiss, which a calibration made with nothing in the loop cannot carry — a genuine difference from the sweep kinds, where the calibration stamp is the floor's first source. The second is the null-run reference (§3.4's control capture, of this kind, under the same calibration — the join is the app's: the newest loop-only Compression record whose embedded calibration snapshot is byte-equal to this record's): the loop's own added-harmonic AMPLITUDE at this step's drive,

$$F_{0,i} = \frac{\mathcal{A}(x_i)}{a_{1,i}}, \qquad \mathcal{A}(x) = 10^{\,\mathrm{interp}_x\left(20\log_{10}(\mathrm{thd}\cdot a_1)\right)/20}, \tag{28}$$

interpolated linearly in dB between the null run's rungs and undefined outside its probed span (an unmeasured floor is never extrapolated). The amplitude, never the null run's THD read across: the loop's contribution is absolute, and the ratio it makes inside another record depends on THAT record's fundamental — reading the ratio across would hand a device with gain a floor too high and a pad one too low by the same amount. The third is a plugin record's operative floor, entering in the same $\sqrt{2}$ shape. The applied floor is the LARGEST of the bounds available — each bounds content that is not this device's distortion, so adding a bound is strictly conservative — and the floor source names a non-default bound only where it actually sets the floor at some point (a reference sitting under this capture's own hiss everywhere is not what the reader is looking at); the null run outranks the stated floor in the label, and in practice they never co-occur.

The classes and the runs. Each point's THD against the applied floor is a floor class: at or under it (to a stated epsilon) the read IS the rig, "below the floor"; above it but within the clear margin, "noise-dominated" (distortion mixed with rig noise); else "clear". The floor runs split the trace into maximal runs by class, each segment between adjacent points wearing the worse endpoint's class and boundary points shared, so the solo and compare charts draw one segmentation; the app draws the three classes solid, dotted and faint, and any segment under its own axis floor as a gap, never as a line at the axis edge. A curve with no stamp (a record from before the stamp existed, or Simulation Mode's) draws with no overlay.

Fundamental presence, summed. The payload stores no per-order levels, so the summed presence reading at a point takes the summed THD as its reference in place of the loudest order: the depth is $20\log_{10}\mathrm{thd}_i$, absent when it reaches chapter one's separation bound. The sum is never smaller than the loudest order, so this reads at most $10\log_{10}(K-1)$ dB more eager than the per-order form — conservative in the refuse-to-quote direction, and nowhere near the empty band that separates real devices from fundamental-absent ones. The bar is borrowed from another instrument: it was measured for the sweep analyzer, and §5.7 measures the one-bin estimator's

own separation, which is far larger. The majority over the points gates the knee, the cleanup reading, the floor label and every quoted THD figure, on the page, in the summaries and in the reader.

The summaries. The generated summaries take four features from the curve — the knee's level with a confidence that caps low for a bound, a sharpness from how many input dB the slope takes to fall between two thresholds, the mean slope over the top of the range, and the cleanup as the same three states — each with literals the parameter table's notes name; they are a later chapter's subject. Two app-level rules the kernel does not own are stated here in one sentence each: a hardware record with a measured volts factor draws its input axis and quotes its knee and cleanup in volts peak at the pedal's input (a plugin record's axis is digital full scale, a different physical quantity, and no cross-domain gap is ever a single number); and three or more records open a symmetric overlay with one chart per input-axis class and no pairwise metric.

5.5 Parameters

Every value below is read from the manifest, which read it from the shipped code: the grids delivered by the shipped generator, the tone's counts by the shipped initializer at its own defaults, the analyzer's order count by measuring a phantom through the default call, and the three default-argument thresholds — the post-knee slope threshold, the cleanup threshold and the at-floor epsilon — by bisecting the shipped functions. The fit's candidate count, ridge, no-break factor, band and decimals bar, the tone's minimum cycle count, the Nyquist guard, the estimator's minimum window, the crossing guard and the summaries' constants are literals the manifest carries as rule strings; the assembly's two silences are the oracle's named statics, mirroring the app's literals.

Parameter	Value
Probe grid	
hardware window (floor / ceiling / steps)	-70 / -12 dBFS / 36
plugin window (floor / ceiling / steps)	-80 / 0 dBFS / 48
hardware / plugin lattice step	1.657 / 1.702 dB
fine window spacing / point cap	0.5 dB / 24
fine-window example (−40 to −34 dBFS): delivered points	48 (the uniform 36 plus 12)
per-step gate margin	6 dB over the noise stamp
Stepped tone	
probe note / sample rate (the initializer's defaults)	220 Hz / 96 000 Hz
settle / measure / ramp durations (default-argument literals)	0.15 / 0.25 / 0.02 s
delivered: settle / measure / ramp samples at the defaults	14 400 / 24 000 / 1 920
measurement window (cycles of the note; minimum)	55 (4)
samples per step / the default grid's tone	38 400 / 1 382 400
Assembly (a mirror of the app's)	
pre-roll / tail silence	0.2 s / 0.2 s
pre-roll / tail / whole stimulus at f_s	19 200 / 19 200 / 1 420 800 samples (14.8 s)
Simulation Mode, interactive: steps / settle / measure / rate	24 / 0.05 s / 0.1 s / 48 000 Hz
Simulation Mode, full: steps / settle / measure / rate	36 / 0.15 s / 0.25 s / 96 000 Hz
Analyzer	

Parameter	Value
harmonic count K (the default, read by measuring)	9
integration band	20 Hz – 20 000 Hz
Nyquist guard (a literal)	$0.95 \cdot f_s/2$
estimator: minimum window (a literal)	8 samples
Knee	
post-knee slope threshold (bisected from the shipped fit)	0.75 dB/dB
pre-knee unity tolerance / the unity guard's bar	0.1 / 0.9 dB/dB
minimum points / candidates / ridge (literals)	4 / 1025 / 1e-9
no-break factor / resolution band / decimals bar (literals)	0.5 · linSSE / 1.05 / 0.5 dB
Cleanup	
threshold (bisected from the shipped function)	0.01 (1 %)
crossing interpolation guard (a literal)	1e-12
Floor	
clear margin	6 dB over the bound
at-floor epsilon (bisected from the shipped helper)	0.1 dB
Fundamental presence	
separation bound	chapter one's, by reference
the summed reference's eagerness, $10 \log_{10}(K-1)$	9.031 dB
Oracle (`analysisdump compression-synth`)	
ladder quadrature points	262 144

5.6 The result

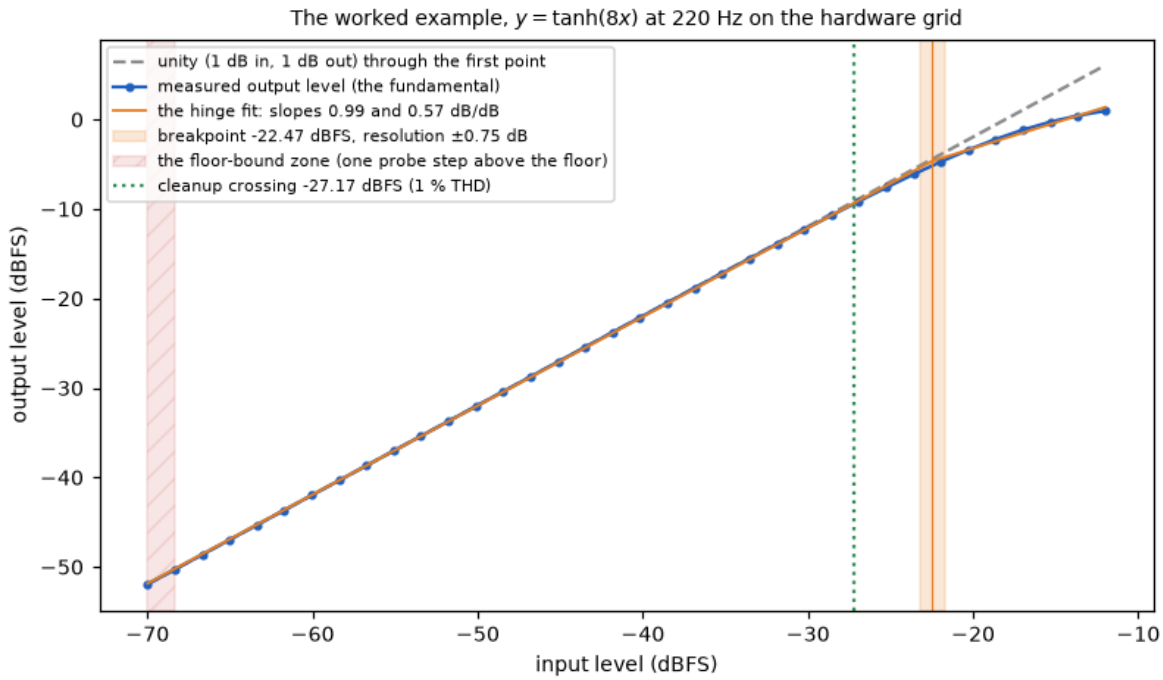


Figure 13: The worked example's curve: $y = \tanh(8x)$ at the default note on the hardware grid, output level against input level with the unity line through the first point, the hinge fit with its two slopes, the breakpoint and its resolution band, the floor-bound zone hatched at the quiet end, and the cleanup crossing.

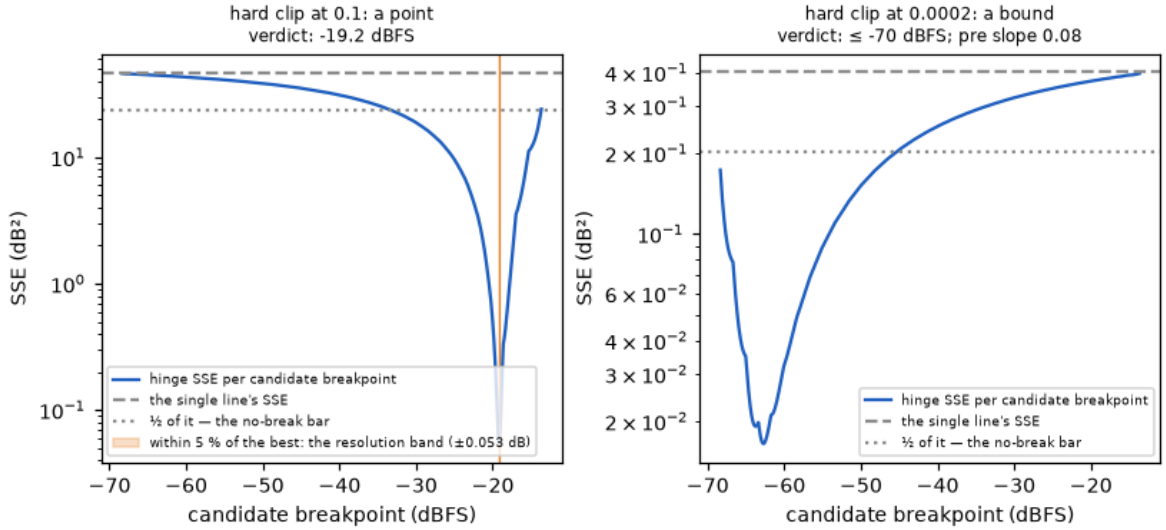


Figure 14: The hinge fit's error over the candidate breakpoints, for a hard clipper whose threshold sits inside the window (left: a sharp minimum, the five-percent band a width, a point) and for one whose threshold sits under the floor (right: the profile falls toward the floor with no anchored minimum, and the verdict is a bound at the floor by the unity guard).

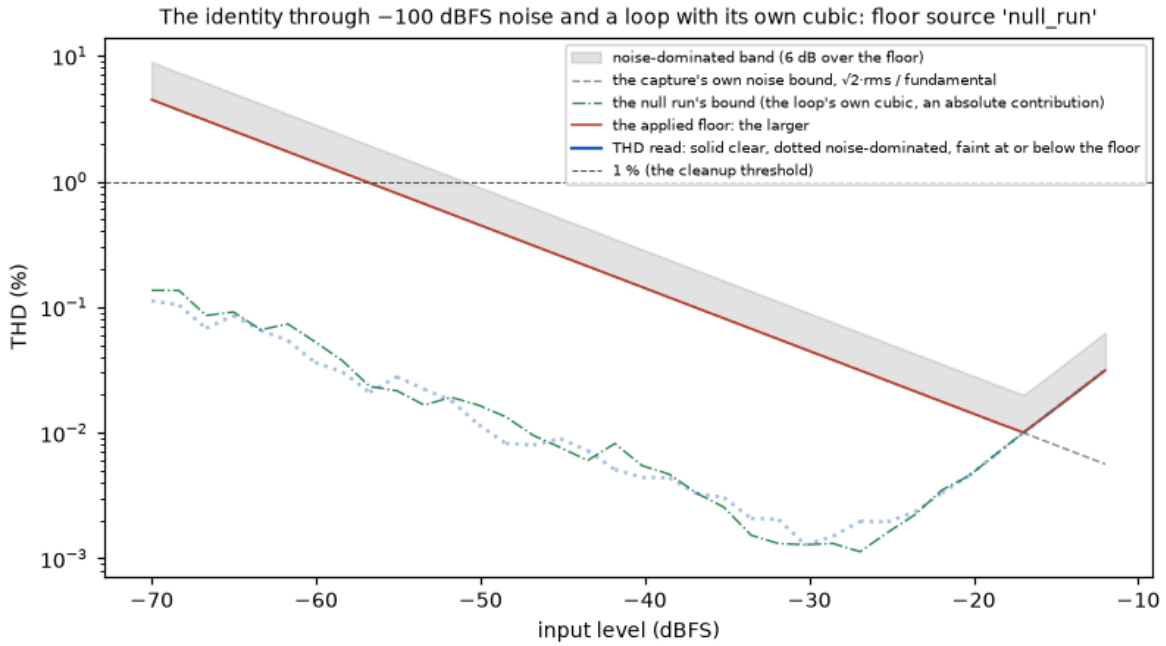


Figure 15: THD against level for the identity through seeded noise and a loop with its own cubic: the capture's noise bound of (27), rising toward the quiet end, the null run's bound (the loop's own distortion, an absolute contribution, defined inside its probed span only), the applied floor as the larger, the noise-dominated band above it, and the trace in the chart's three styles — here at or below the floor everywhere, since the identity adds nothing of its own.

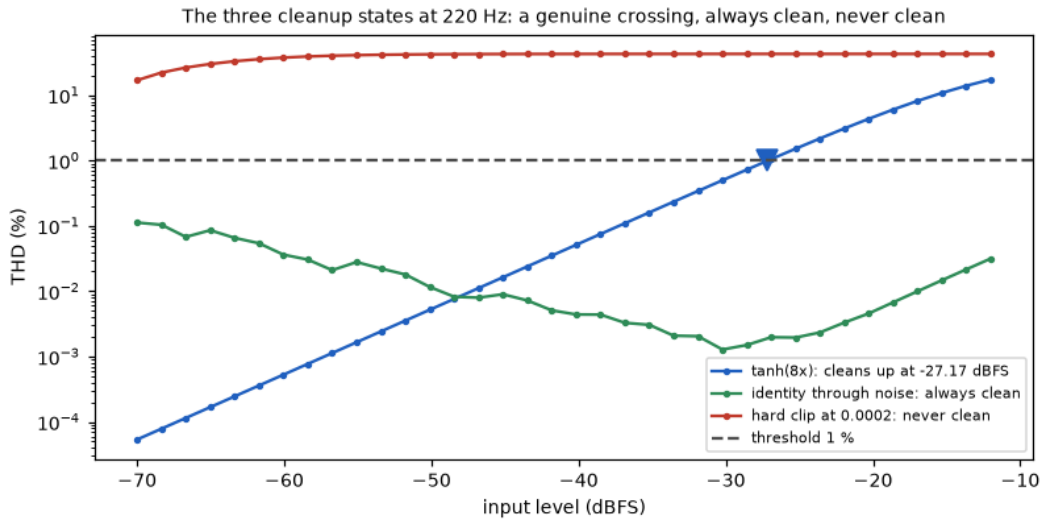


Figure 16: The three cleanup states: a genuine crossing interpolated between the two steps that straddle the threshold, a curve under the threshold at every step (always clean), and a clipper whose threshold sits under the grid's floor (never clean).

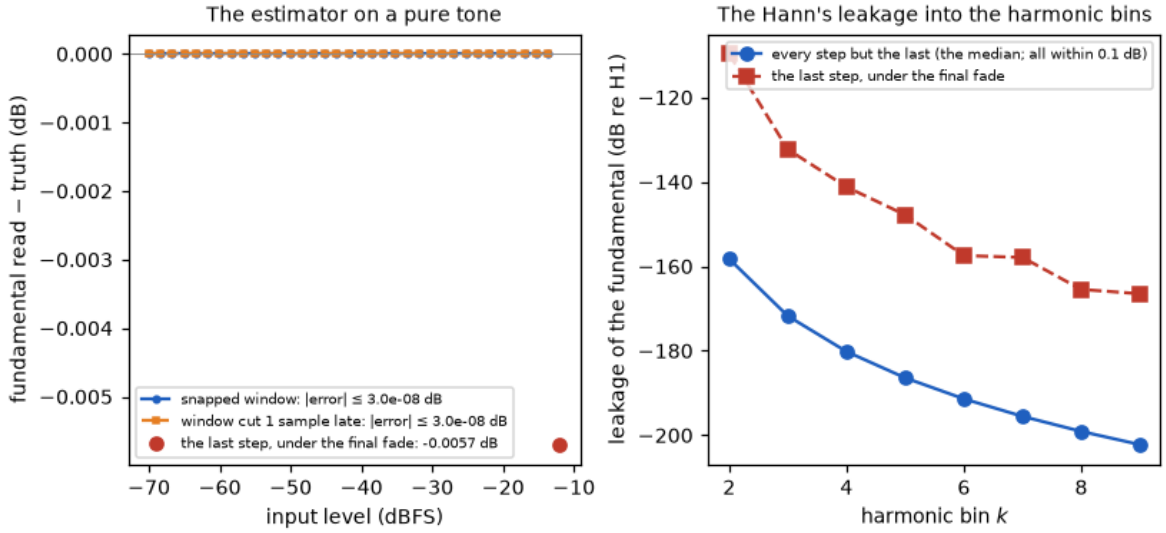


Figure 17: The estimator on a pure tone (left): the fundamental's read against its truth per step on the snapped window and on a window cut one sample late — both at rounding — and the last step, which holds the stimulus's final fade; and the Hann's leakage of the fundamental into every harmonic bin (right), by bin distance, on the snapped window and under the fade.

The app's own rendering of the same measurement kind, from the user guide's figure pipeline (Figure 18).



Figure 18: The Compression view as the app draws it, from the guide's simulated fixture: the measured output level following the dashed unity reference at low levels then flattening, the knee and cleanup tiles, and the companion THD trace below with its gap under the chart's own floor.

5.7 Worked example and parity

The parity test runs 24 cases through both implementations and two device cases again as their phantom capture twins (2 twins). The Python imports chapter one's device family and noise generator and the transfer chapter's filters, and owns the rest: this instrument has no sweep. Every case is a known device through a known LOOP — a delay, one seeded noise stream referred to the interface's input, a chain gain applied after both, and optionally the loop's own cubic, the converter's distortion — inside the app's assembly, and beside every step the oracle prints its truth: the memoryless device's closed-form ladder at that step's drive (chapter one's quadrature, one pass for all orders, the fundamental's own compression included, the loop's cubic composed where it is set), scaled by a pre-filter's magnitude at f_0, a post-filter's at kf_0 and the chain gain; the truth THD as (25)'s order rule applied to that ladder; and the read-time half as the shipped predicates run on a curve built from the truth points — the definition applied to the ladder, never a second derivation. A knee has no analytic truth beyond that: a hinge fit is a summary of a curve, not a property of the device.

The instrument every truth comparison rests on is the estimator's own leakage: the pure phantom [1], whose truth is zero at every order above the first, read on the tone's own geometry. The symmetric Hann leaks a tone into the other bins by bin distance — -158.2 dB re the tone at one bin, -171.8 at two, -202.2 at eight — level-invariant to rounding:

bin distance	1	2	3	4	5	6	7	8
leakage (dB re the tone)	-158.2	-171.8	-180.2	-186.4	-191.4	-195.5	-199.1	-202.2

A content order can read up to its truth plus what every OTHER present order leaks into its bin, $20\log_{10}(1+\Lambda_k/a_k)$ with Λ_k the sum by distance; that allowance is subtracted before the bar is applied, and what remains is the method's own error (measured: the read sits AT the allowance, to two percent). Under the final fade the one-bin leakage rises to -109.7 dB, which is why the last step is excluded from every comparison and why the fade is a stated limit. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. The test asserts:

- Swift against Python: every header quantity and every measured column within 1e-07 dB (measured maximum 3.1e-08 on this run) on a capture with no filter in it and within 1e-06 (measured 1.6e-08) through a biquad; the two quadratures' truth columns within 1e-06 (measured 3.3e-07); every discrete reading — the gate's verdict per step, every class, every run, the knee's state and bound, its label, the cleanup state, the floor label with and without the null run, the presence verdict — identical; dust (a harmonic whose truth sits more than 140 dB under the fundamental) is stated and not compared;

- the shipped read against the ladder on every noiseless case, every step but the last: the fundamental within 1e-06 dB (measured 9e-07); the content harmonics after the allowance within 0.0002 for the smooth devices (measured residue 6.9e-05), 0.0002 for the hard clipper (8.2e-05 — at the default note and rate no alias image of a per-sample device lands on a harmonic bin, since the rate is not an integer multiple of the note, so the kinked device reads like a smooth one here; its twin 5.8e-05) and 0.005 through a filter (0.0018); the THD after the allowance summed over every bin the analyzer sums within 0.0001 (7.8e-06); the last step's fundamental LOW by the fade, between -0.0057 and -0.0017 dB (the bar 0.02 — a compressing device flattens the fade);

- the knee against the truth curve's knee within 0.1 dB (measured 0.0534 — every point knee on the hardware grid lands on the same candidate, and the two that differ do so by one candidate step), and the cleanup crossing within 0.001 (0.00034);

- the stamp against the injected noise within 0.02 dB (measured 0.00744 — the pre-roll estimate's own scatter); the same device and noise through a loop 12 dB louder: the knee's input level moved by $\leq$ 1e-10, every class equal, every output level moved by 12 dB with a departure of $\leq$ 1e-10 (bar 1e-09); a capture 100 samples late with the analyzer told the pre-roll plus that: every read against the undelayed capture within $\leq$ 1e-10 (bar 1e-09); a window cut 1 sample late: the fundamental moved by 3e-08 dB (bar 1e-06) — the Hann's taper makes a one-sample mis-cut invisible;

- Part B on the oracle's loops: the identity through noise reading every point at or below the noise bound (the bathtub), always clean, no knee; the null run through a linear loop labeled capture noise (its bound sits under this capture's stamp everywhere) and through a loop with its own cubic labeled null run, its bound nil outside its span; the four verdict legs on their fixtures; the fine window keeping the local probe step; the plugin window's ceiling at exactly 0 dBFS; the stated floor governing; the device with no fundamental refused everywhere; and the one-bin estimator's own separation on a phantom with no fundamental, 158.2 dB against the borrowed 76 dB bar (pinned to stand over it by 60 dB).

Part B's parity is established: on every read-time case — the identity through noise with and without a null run, the null run through a loop with its own distortion, the exact sub-unity lines,

the clippers at and under the floor, the fine window, the plugin window, the stated floor, and the device with no fundamental — the two implementations agree on every bound, every class, every run, every knee field, every cleanup state and level, every floor label and every verdict.

The device cases, each with its content-harmonic count, its Swift-vs-Python figure, its fundamental's error, its worst raw harmonic error and where, its residue after the allowance, its THD residue and its fade:

case	content	S−P (dB)	H_1 (dB)	raw (dB)	at	residue (dB)	THD residue	fade (dB)
a-tanh8	71	2e-10	3e-08	0.042	step 0 h3	5.8e-05	≤ 1e-10	-0.0031
a-tanh8-twin	71	2.6e-10	3e-08	0.042	step 0 h3	5.8e-05	≤ 1e-10	-0.0057
a-tanh2	47	1e-09	3e-08	0.212	step 3 h3	6.9e-05	≤ 1e-10	-0.0054
c-hardclip	16	≤ 1e-10	9e-07	8.22e-05	step 34 h9	8.2e-05	7.8e-06	-0.0017
c-hardclip-twin	16	≤ 1e-10	3e-08	9.73e-07	step 33 h5	1.2e-07	3.7e-09	-0.0057
l-tanh-prehp	71	1.6e-08	2.6e-08	0.0416	step 0 h3	1.2e-08	≤ 1e-10	-0.00295
l-poly-postlp	63	1.5e-08	2.9e-08	0.21	step 7 h3	0.0018	≤ 1e-10	-0.00558

Every case's read-time verdicts on the Swift side — the knee the page would print, the cleanup, the floor label — with the two implementations' agreement:

case	knee	cleanup	floor source	S−P (dB)
a-tanh8	-22 dBFS	cleans up at -27.17 dBFS	none	2e-10
a-tanh8-twin	-22 dBFS	cleans up at -27.17 dBFS	none	2.6e-10
a-tanh2	none	cleans up at -15.13 dBFS	none	1e-09
a-linear	none	always clean	none	≤ 1e-10
b-identity-noise	none	always clean	capture noise	≤ 1e-10
b-identity-null	none	always clean	capture noise	≤ 1e-10
b-identity-null-loop	none	always clean	null run	≤ 1e-10
b-tanh8-null-loop	-22 dBFS	cleans up at -26.88 dBFS	null run	≤ 1e-10
c-hardclip	-19.2 dBFS	cleans up at -20.03 dBFS	none	≤ 1e-10
c-hardclip-twin	-19.2 dBFS	cleans up at -20.03 dBFS	none	≤ 1e-10
d-subunity-break	≤ -70 dBFS	cleans up at -30.18 dBFS	none	≤ 1e-10
d-subunity-line	none	cleans up at -50 dBFS	none	≤ 1e-10
d-hardclip-zone	≤ -70 dBFS	cleans up at -68.06 dBFS	none	≤ 1e-10
d-hardclip-under	≤ -70 dBFS	never clean	none	≤ 1e-10
e-fine	-19.2 dBFS	cleans up at -19.77 dBFS	none	≤ 1e-10
f-plugin	-18.5 dBFS	cleans up at -20.14 dBFS	none	≤ 1e-10
g-gain0	-22 dBFS	cleans up at -27.17 dBFS	capture noise	≤ 1e-10
g-gain12	-22 dBFS	cleans up at -27.17 dBFS	capture noise	1.3e-10
h-latency	-22 dBFS	cleans up at -27.17 dBFS	none	2e-10
h-miscut	none	always clean	none	≤ 1e-10
i-operative	-22 dBFS	cleans up at -27.17 dBFS	stated floor	≤ 1e-10
j-absent	not_stated	none	none	≤ 1e-10
j-separation	not_stated	none	none	3.1e-08
k-simulation	-21.2 dBFS	cleans up at -27.14 dBFS	none	4.6e-10
l-tanh-prehp	-22 dBFS	cleans up at -27.17 dBFS	none	1.6e-08
l-poly-postlp	none	cleans up at -20 dBFS	none	1.5e-08

The worked example, $y = \tanh(8x)$ noiseless on the hardware grid: 71 content harmonics compared, the fundamental within 3e-08 dB, the worst raw harmonic error 0.042 dB at step 0 h3 and 5.8e-05 dB after the allowance (at step 28 h9 (74 dB clear)), a THD residue of ≤ 1e-10 dB, the last step -0.0031 dB under its truth. Its knee is a point at -22.469 dBFS (resolution ±0.748 dB, displayed as "-22 dBFS"), slopes 0.9936 and 0.5749 dB/dB, against the truth curve's -22.469 (a difference of ≤ 1e-10); its cleanup crossing -27.175 dBFS against the truth curve's -27.175 (a difference of 1.3e-06); its summaries read a sharpness of 0.586 and a top-of-range slope of 0.478 dB/dB. The hard clipper at a threshold of 0.1: clipping begins at -20 dBFS and the fitted breakpoint sits at -19.16 dBFS, 0.842 dB above the onset, with a post-knee slope of 0.157 and a resolution of ±0.0534 dB — a MEASURED relation between a fit's corner and a clipper's onset, not a claim; with the fine window over that knee the grid delivers 52 points, the knee reads -19.21 dBFS at ±0.0534 (the resolution unchanged — on both grids the five-percent band is narrower than one candidate step, so the candidate spacing is what the fit resolves) and moves by 0.0534

dB. The exact sub-unity line with a break reads "≤ -70 dBFS" with fitted slopes 0.85 and 0.2 (the unity guard). The device with no fundamental — a phantom whose fundamental sits 80 dB under its second harmonic — reads absent on 0.833 of its points at a median depth of 79.97 dB and no knee, no cleanup and no floor label. Simulation Mode's assembly on the worked example (interactive quality: fewer steps, shorter windows, no pre-roll) reads its knee at -21.19 dBFS at ±0.388, 1.28 dB from the app assembly's — what interactive quality costs on this device.

5.8 Limits and failure modes

- **The last step is measured under the stimulus's own fade.** The tone's final fade covers its last ramp length, and the last step's measurement window ends where the tone ends, so that window holds the fade: the last step's reads sit low (-0.0057 to -0.0017 dB on the parity's devices) and the one-bin leakage there rises from -158.2 to -109.7 dB. A stimulus fact, not a device's — moving the fade past the window is a capture-time change to the corpus's stimulus and is reported, not made.
- **The no-break test has no scale, and the fade can decide it.** The first verdict leg compares two residuals as a RATIO; on a curve that is one exact sub-unity line the only departure from a line is the last step under the fade, the hinge absorbs that one point, the ratio reads a break, and the second leg then reads no knee — where the same exact line with no fade reads the bound at the floor. Measured: the exact line reads "none" on both implementations while its truth curve reads "≤ -70 dBFS" with a pre-knee slope of 0.85; the kernel's own fixture for that leg has no fade and passes. The exact line with a break is unaffected (the unity guard fires on its fitted pre slope). A measured device rarely draws an exact line, and any real curvature overwhelms one bent point; the leg is stated as decided by structure the test cannot see, and the two mechanisms above are one finding.
- **Skipped steps are silent.** A step whose window runs past the capture is dropped with no count, unlike the gate's; a short capture yields a shorter curve, and whether the omission deserves a stamp is a stored-format question, reported here.
- **The assembly lives in four places.** The record's pre-roll, tail, stamp and cut are the app's literals, copied in the command-line tool and written differently twice more; the oracle mirrors the app's and says so, and a shipped plan is the filed follow-up. Until it lands, the mirror is the one place this chapter's pipeline is not the shipped code called.
- **A soft onset is not disclosed.** The unity guard tells a clean anchor from a missing one; it does not tell a wide soft onset from a corner. A drive-max soft-knee overdrive record in the library reads a unity bottom decade and a wide soft onset above it, and its point knee is accurate as far as the fit goes; a transition-width disclosure would be a new mechanism, not a tolerance.
- **The borrowed bar.** The presence separation was measured for the sweep analyzer; the one-bin estimator separates a fundamental from its harmonics by 158.2 dB on a snapped window, so the bar is conservative here by a wide margin, never optimistic.
- **One note.** The curve is one pitch; a device with a tone network ahead of its clipper compresses differently at another, and the Gain Map (§6) is the cross-check across notes.
- **The quiet end is the rig's.** Below the noise bound the read is the rig, and the classes say so per point; a low-level upturn in THD is the bathtub, not the device, unless it stands clear of the band.
- **Aliasing in a simulated device.** A device applied per sample folds its content above Nyquist onto the reads; at the default note and rate no image lands on a harmonic bin, and

the parity's twins are the alias-free control. At another note or rate a kinked device's high orders would not be the oracle for that bin.

6 Gain Map

6.1 Significance and limitations

A gain map is the Harmonic Distortion measurement of §3 repeated at each drive level of a level lattice — one sweep per level, each level analyzed on its own — and summarized on one grid: for every level i and every column j of an export grid of fundamentals, the total harmonic distortion over the audible band and the level of every stored harmonic order, each a cell. It answers what one sweep cannot: how the harmonic content moves with how hard the device is driven — whether it cleans up, where, and whether its recipe holds across the range. At read time the record states a peak tile (the loudest cell that is a measurement, or a bound), a cleanup reading per column and one tile at the Compression probe's note, and a fundamental presence verdict over the map.

What it cannot tell you, stated once: (a) anything about the device's output level or gain — the cells are ratios to each level's own fundamental, and the drive axis is the signal driving the circuit, never a knob on it (to map a knob is to measure one map per position); (b) anything between two levels — the rows are separate captures seconds apart, and a circuit that drifts or sags between them shows the drift as row-to-row wobble; (c) whether a cell near its floor is the device or the rig — below the composed margin a cell is quoted as a bound; (d) anything at the two edge mask columns, which are the sweep's own fade ramps; (e) anything a truncated column lost above the top of the band; and (f) anything about what the distortion sounds like.

6.2 The stimulus

The lattice is N levels evenly spaced in dB from a floor to a ceiling inclusive,

$$\ell_i = \ell_{\min} + (\ell_{\max} - \ell_{\min})\,\frac{i}{N-1}, \qquad A_i = 10^{\ell_i/20}, \qquad i = 0 \ldots N-1, \tag{29}$$

a count of one being the ceiling alone. Two drive windows exist. The hardware window is anchored to gain staging: with the interface's input staged so that pedal peaks land in the staging target at the loudest planned drive, the ceiling is the loudest drive that cannot clip the converter, and the floor sits a fixed travel below it (a step of 5.143 dB on the default lattice). The plugin window's ceiling is digital full scale — the float render path has no converter to protect — and its floor is lower, where high-gain preamp knees were measured to sit (a step of 8.571 dB). The parameter table gives both.

Each level's stimulus is the synchronized sweep of §3.2 at that level's amplitude — the same band, the same requested length, the same fades — between a pre-roll of silence (the per-level noise region) and a tail of silence:

$$s_i[n] = [\underbrace{0, \ldots, 0}_{\lfloor t_{\text{pre}} f_s \rfloor},\ x_i[0], \ldots, x_i[S-1],\ \underbrace{0, \ldots, 0}_{\lfloor t_{\text{tail}} f_s \rfloor}], \tag{30}$$

with S the sweep's sample count. The sweep length is the measure sheet's pick from a short list and is *not stored* in the record; the consequence is a floor derivation that must take the worst case over the list (§6.4). The default length is shorter than the fingerprint sweep's, so every extraction window is shorter and the analyzer separates orders less well at the band's low edge (§6.7 measures it).

6.3 The pipeline

Given the N captured outputs $y_i[n]$ and the loop's integer latency ℓ (the calibration's own measured figure on the hardware path; zero on the plugin path, which measures and removes its latency per capture):

1. **The aligned cut.** From each capture the sweep and its tail are cut at the pre-roll plus the latency, $y_i[\lfloor t_{\text{pre}} f_s \rfloor + \ell \;\ldots\; \lfloor t_{\text{pre}} f_s \rfloor + \ell + S + \lfloor t_{\text{tail}} f_s \rfloor]$, clipped to the capture's end — one arithmetic for every caller.

2. **The per-level analysis.** Each aligned capture goes through the analysis of §3.3 with the level's own sweep and the journey's configuration: the same harmonic count, a response DFT half the fingerprint analyzer's length (twice its bin width), the same window cap. The levels are analyzed in parallel — the analyses are pure — and the entries sort by ascending amplitude. Nothing is compensated: the loop's gain and response stay in every order.

3. **The summary grid.** On the export grid $f_j = f_1(f_2/f_1)^{j/(n-1)}$, $j = 0 \ldots n-1$ — n fundamentals log-spaced from the sweep's start to its end *inclusive*, so the first and last columns sit at the sweep's edges — each level's THD over the audible band B and each stored order's level:

$$\text{thd}[i][j] = \frac{\sqrt{\sum_{k\geq 2,\, kf_j \in B,\, \text{valid}_i(k,f_j)} |H_k^{(i)}(kf_j)|^2}}{|H_1^{(i)}(f_j)|}, \qquad h[k][i][j] = 20\log_{10}|H_k^{(i)}(kf_j)|, \quad (31)$$

   nil where the fundamental reads zero or the harmonic sits at or above Nyquist; valid_i is the level's own validity bound — the sweep's alias-image bound and nothing else, since the journey is uncompensated and has no calibrated band. Because the cells are uncompensated, a cell's delivered fundamental, $\ell_i + h[1][i][j]$, is the level the note actually reached at the converter in absolute dBFS; the floor's drive term reads it.

4. **The models.** Beside the grid the record carries one audition model per level — a Hammerstein fit of that level's analysis, bounded in level by its own amplitude. The Hear-it panels consume them and nothing measured does; a render between two levels blends two fits and the quieter one clamps, inventing a bracket-shaped clipping the device does not have, measured at −25 dB on this lattice — a property of what the models *render*, stated here and measured in a later chapter.

5. **Storage.** The grid, the levels, the orders, the models, the sample rate and the integration band. The last two are additive fields: a legacy record without them derives no validity, no fade, no masks and no analyzer residue, and its cells summed everything below Nyquist.

6.4 Read-time predicates on the result

Everything below is computed from the frozen grid when a record is read.

The validity fade and the truncated columns. The summary's validity is the grid's own span with the alias-image bound: a harmonic's cells past $f_j \leq f_s/(2k+1)$ are faded on display. A truncated column is one where some order $k \geq 2$ has kf_j inside the stored band and fails the bound: its THD cells under-read their declared band, and the column is flagged and excluded from the tiles, never corrected.

The edge masks. The two columns within a quarter tone of the grid's edges — $f_j \leq f_0\rho$ and $f_j \geq f_{n-1}/\rho$, ρ the guard ratio — are the sweep's fade-in and fade-out ramps: excited under a ramp, they describe the measurement and not the device (a jumpered wire reproduces them). They are derived from the grid alone, faded, excluded from the tiles and from the presence vote, and never touch the color scale.

Fundamental presence. Per cell, the reference is the loudest order *stored* for that cell, the fundamental included; the cell is absent when the fundamental sits at or beyond the analyzer's separation bound of §3.4 below it, present otherwise, undetermined without a fundamental read; no rig floor enters. The map's verdict is the majority over every cell outside the two edge columns, and it outranks both peak states. The vote is a majority for a reason the worked example shows: a cell whose fundamental bin holds the loop's noise, or — beside the low edge — the second harmonic's own leakage into the fundamental's window, reads present on a device that has no fundamental at all.

The floor, composed from three sources. A cell's THD is quoted as a measurement only where it stands clear of *every* bound the rig and the analyzer put under it; the applicable floor is the largest, and the largest floor is the smallest margin.

The first source is the calibration stamp: the embedded calibration's THD-versus-frequency curve of the loopback, one read per point at log-spaced fundamentals, measured once at the calibration's own drive. It is read through the median read: each point replaced by the running median of its neighbors' dB over an odd window before the interpolation on the $(\log f, \mathrm{dB})$ plane. The window is a fixed count of points on the stamp's own grid (± 0.25 octave, 0.5 octave in all — about 1.3 of the map's 0.39-octave columns, i.e. about as wide as one column) — wide enough that a single-read trough cannot set the floor. The stamp is a ratio re the calibration's delivered fundamental, while the cells are uncompensated, so the rig's absolute noise is the calibration's level plus its measured chain gain, and a cell's floor ratio is that against the cell's delivered fundamental:

$$F_{\mathrm{rig}}[i][j] = \mathrm{stamp}_{\mathrm{med}}(f_j) + \max\bigl(0,\ (\ell_{\mathrm{cal}} + G) - (\ell_i + h[1][i][j])\bigr), \tag{32}$$

the drive term. Quieter than the calibration the ratio rises by the difference, because the noise is absolute; louder, it stays at the stamp's own ratio — never extrapolated past what the calibration measured, because the rig's own distortion rises with drive and the stamp measured their sum at one drive (the bare loop at the loudest level read a few thousandths of a percent at 6–16 dB "clear" of a floor extrapolated 14 dB below anything the calibration saw). The calibration source contributes only with both a usable stamp and a chain gain; omitting the gain left every margin one input-knob setting optimistic.

The second source is a plugin record's operative floor — the larger of its silence render and its A/A residual, an rms re full scale — entering as $20\log_{10}(\sqrt{2}\,\mathrm{rms}) - (\ell_i + h[1][i][j])$, an absolute bound with no chain-gain term, because the grid is already absolute. It governs where the render-path calibration is numerically silent. The larger of the two ratios is the rig-noise floor; nil when neither source is usable, since a floor known to be missing a term is worse than none.

The third source is the analyzer residue, and it is what a noise floor cannot see. The plan's own sweep, handed back as its own capture with the tail's silence — nothing in the loop — goes through the shipped analyzer at the journey's configuration, and its THD is read per column of the record's grid: the THD of nothing. One column above the sweep's start the fade-in ramp leaks *coherently* into the second-harmonic window: on the shipped grid the residue at 26.2 Hz reads -70.11 dB on the default sweep, -46.32 on the short one and -86.67 on the long one, level-invariant — a ratio — and a column up it is -103.6 dB, under every real stamp. The sweep length is not stored, so the record takes the worst (loudest) residue per column over the offered lengths, -46.32 dB at that column; the residue is computed at read time from the record's own geometry and cached per grid:

j	f_j (Hz)	2 s	5 s	10 s	worst
0	20	-21.63	-27.74	-34.13	-21.63
1	26.2	-46.32	-70.11	-86.67	-46.32
2	34.33	-74.85	-103.6	-117	-74.85
3	44.99	-92.87	-117.2	-129.3	-92.87
4	58.94	-104	-127.2	-140.3	-104
5	77.23	-114.2	-136.3	-150.4	-114.2
6	101.2	-123	-144.4	-160	-123
7	132.6	-130.2	-152.1	-169	-130.2

The composition is per cell,

$$m[i][j] = \min\Big(20\log_{10}\mathrm{thd}[i][j] - F_{\mathrm{rig}}[i][j],\ 20\log_{10}\mathrm{thd}[i][j] - R[j]\Big), \tag{33}$$

the composed margin; a cell with no rig-noise margin keeps nil, a column with no residue keeps the rig-noise margin.

The peak tile. The truncated and edge columns masked, the cells whose presence reads absent dropped, the map-level verdict first (absent ⇒ no percentage at all, only the median depth). Pass one: a cell is *clear* when it is unmasked, measured, and its composed margin reaches the minimum margin — every unmasked measured cell when there are no margins. Pass two: a clear cell is a corroborated peak when the cell one column either side, the same level one column (0.39 octave, about 4.7 semitones) away, is also clear. The shoulder is asked for in frequency and never in level: a nonlinearity acting at one drive acts on the neighboring notes too and THD moves slowly across one column's 0.39 octave, while across one level step it can move by tens of dB, so a genuine onset peak at the loudest level may have no clearing level neighbor — and a noise pair in adjacent levels of a single-track column is as likely as any other pair (measured: two of five bare-loop nulls paired that way). The tile is the loudest corroborated cell as a peak; else the loudest unmasked cell as a bound, quoted with ≤; else nothing. Without floor information — Simulation Mode, a record with no calibration — every unmasked measured cell counts as clear and corroborated, and the raw maximum is stated as a peak: reported, not changed.

The cleanup reading. A column's $(\ell_i, \mathrm{thd}[i][j])$ pairs, ascending, cells without a THD dropped, read as one of three states: the first point already at or above the threshold is *never clean* (within this map's range); the first i with $\mathrm{thd}_i \geq \tau$ gives *cleans up* at

$$\ell_{i-1} + \frac{\tau - \mathrm{thd}_{i-1}}{\max(\mathrm{thd}_i - \mathrm{thd}_{i-1},\ \epsilon)}\,(\ell_i - \ell_{i-1}), \tag{34}$$

else *always clean* — the endpoint of the measured range is never dressed up as a crossing. The threshold is a default-argument literal in the kernel, read for the manifest by bisecting the shipped initializer; ϵ is a literal guard. The tile reads the column nearest, in log frequency, the Compression probe's own note, so the two records' cleanup figures are about one pitch, and it is gated on the presence verdict. The summaries' journey fallback, used only when the Compression kind supplied no cleanup, reads the column nearest the same note in *linear* frequency with at least three levels, the same interpolation and a fixed confidence (all literals). The page adds one rule outside the kernel: a never-clean tile is scoped to the map's own drive range in its subtitle, and when the pedal's variant-matched Compression record finds a genuine crossing below the map's floor the subtitle points there.

6.5 Parameters

Every value below is read from the manifest, which read it from the shipped code; the residue capture's drive, the summaries' reference note and confidence and level count, the crossing guard,

the calibrator's grid start, point count and cut, and the chain gain's grid are literals the manifest carries as rule strings.

Parameter	Value
Drive windows	
hardware window (floor / ceiling / levels)	-48 / -12 dBFS / 8
plugin window (floor / ceiling / levels)	-60 / 0 dBFS / 8
hardware lattice step	5.143 dB
Plan	
sweep band f_1 – f_2	20 Hz – 10 000 Hz
sweep length (default) / offered	5 s / 2, 5, 10 s (not stored)
actual duration / rate constant L at the default	4.97169 s / 0.8 s
pre-roll / tail silence	0.3 s / 0.5 s
sample rate f_s (the manifest's)	96 000 Hz
sweep samples / pre-roll / tail at f_s	477 282 / 28 800 / 48 000
analyzer: harmonic count K / response DFT M / half-width cap	9 / 32 768 / 0.2 s
export grid columns	24 (20–10 000 Hz inclusive)
Summary	
integration band	20 Hz – 20 000 Hz
edge mask guard ratio (a quarter tone)	1.0293022
masked columns on the default grid (bottom / top / truncated)	0 / 23 / none
Floor	
stamp median window (the Gain Map) / per-harmonic path	7 points / 1 (verbatim)
calibration analyzed sweeps / level	1 / -26 dBFS
calibration sweep (the app's)	20 Hz to round(0.45 f_s), 10 s, 0.5 s before / 1 s after
residue capture drive	-26 dBFS (a literal; the read is a ratio)
Peak tile	
minimum margin	6 dB
Cleanup	
threshold (bisected from the shipped initializer)	0.01 (1 %)
tile column (nearest in log frequency)	220 Hz
summaries' fallback: reference / minimum levels / confidence	220 Hz (linear nearest) / 3 / 0.7 (literals)
crossing interpolation guard	1e-12 (a literal)
Oracle (`analysisdump journey-synth`)	
ladder quadrature points	262 144

6.6 The result

The app's own rendering of the same measurement kind, from the user guide's figure pipeline (Figure 24).

6.7 Worked example and parity

The parity test runs 15 cases through both implementations and two device cases again as their phantom capture twins (2 twins). The Python imports chapter one's reimplementation for the sweep, the device, the phantom, the noise and the analyzer, and adds only what the map adds.

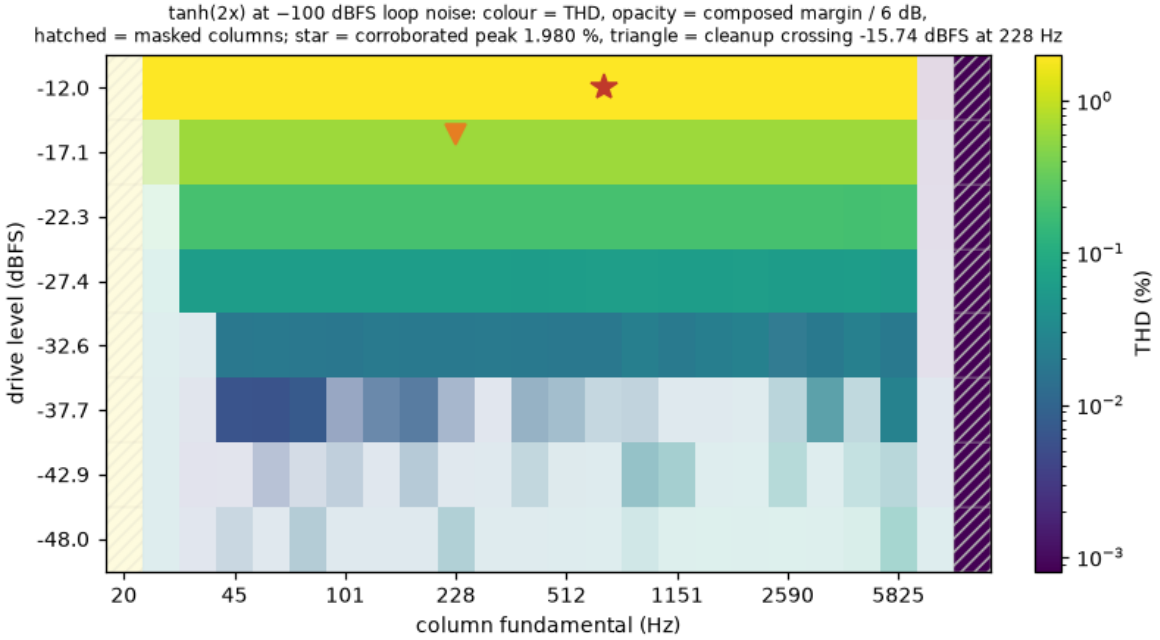


Figure 19: The worked example's map: $y = \tanh(2x)$ on the hardware lattice through a loop with seeded noise, every cell colored by its THD and drawn at an opacity that follows its composed margin, the two edge columns hatched, the corroborated peak starred and the cleanup crossing marked at the tile's column.

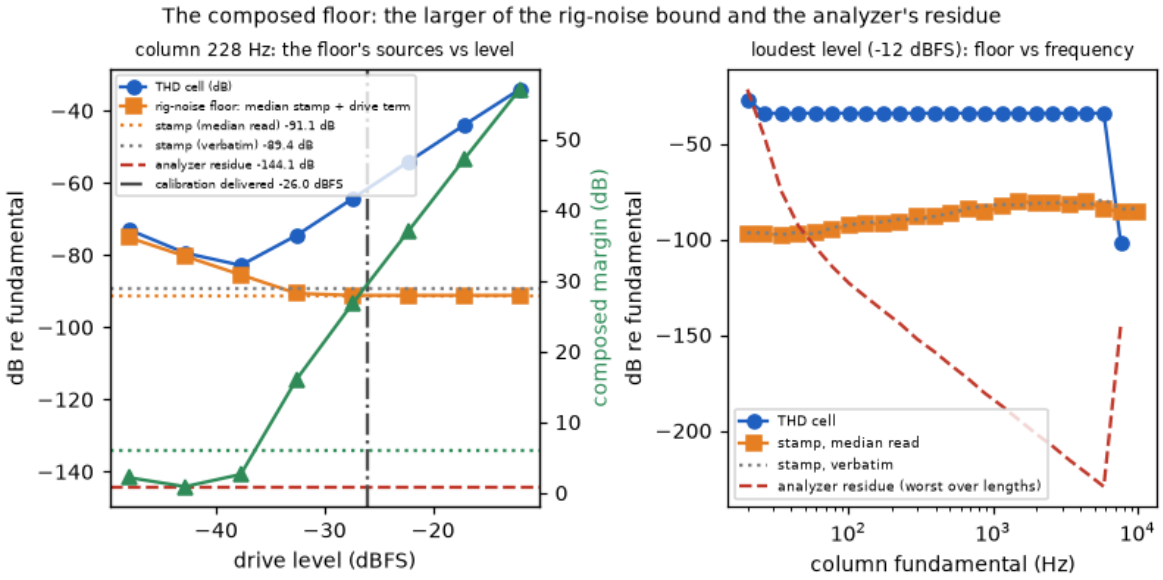


Figure 20: The composed floor at one column against level (left): the stamp read verbatim and through the median, the drive term rising below the calibration's delivered level and clamped above it, the analyzer's residue, and the composed margin against the minimum; and at the loudest level against frequency (right).

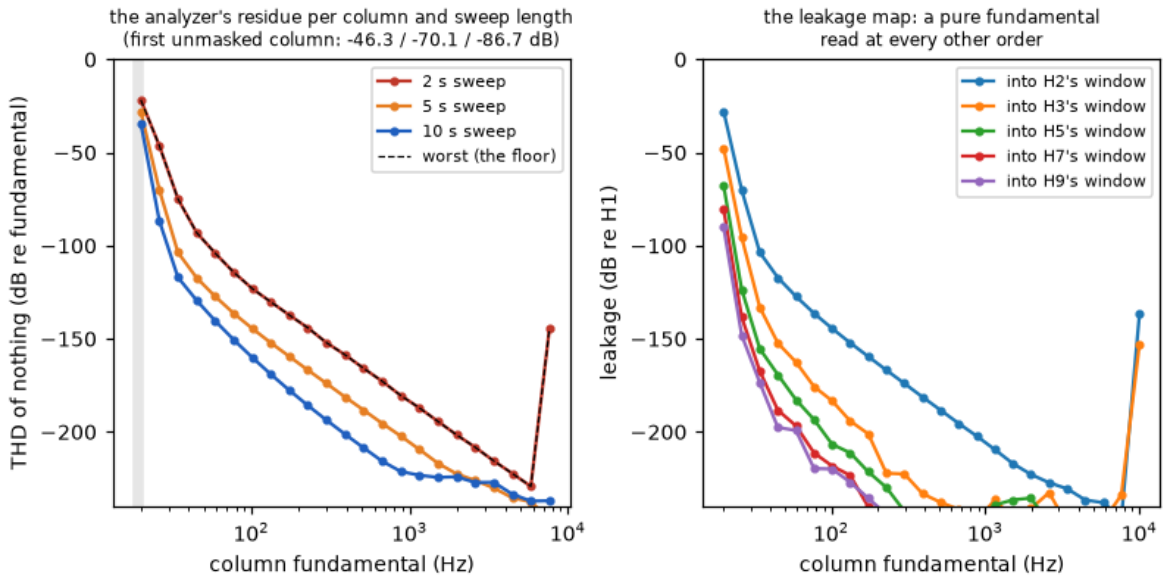


Figure 21: The analyzer's residue per column at the three offered sweep lengths (left) — the fade-in leak at the first unmasked column and the worst envelope the record takes — and the leakage map (right): a pure fundamental read at every other order, the parity's own allowance.

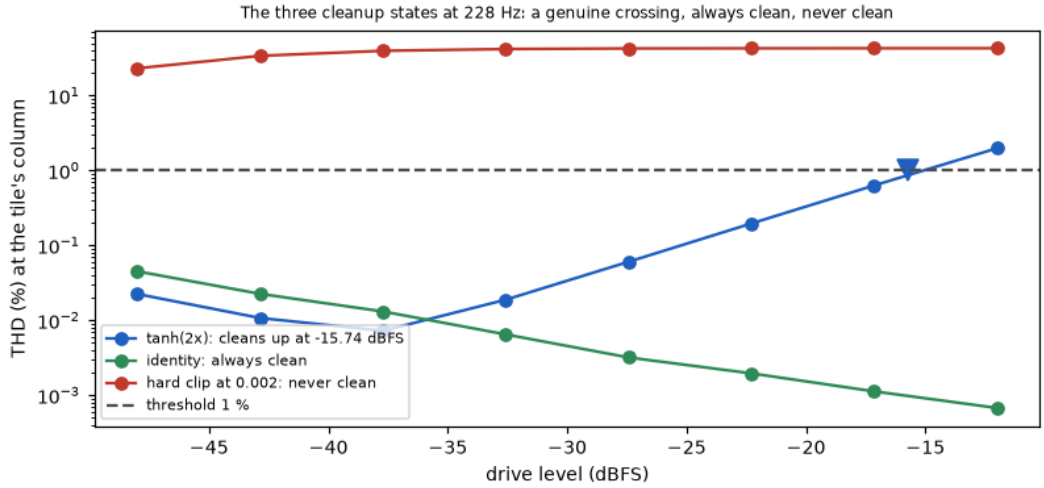


Figure 22: THD against level at the tile's column for the three cleanup states: a genuine crossing interpolated between two levels, a column under the threshold throughout, and a clipper whose threshold sits under the lattice's floor.

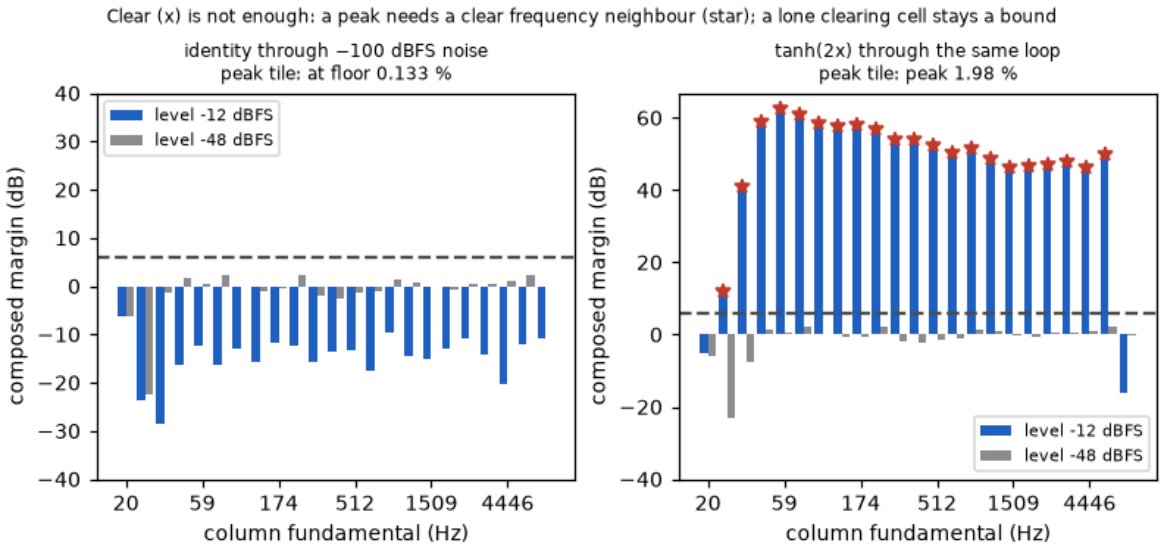


Figure 23: The corroboration rule: composed margins per column at the loudest and the quietest level for the identity through noise (left, nothing clears and the tile is a bound) and for the soft clipper through the same loop (right, every clear cell has a clear neighbor and the tile is a peak).

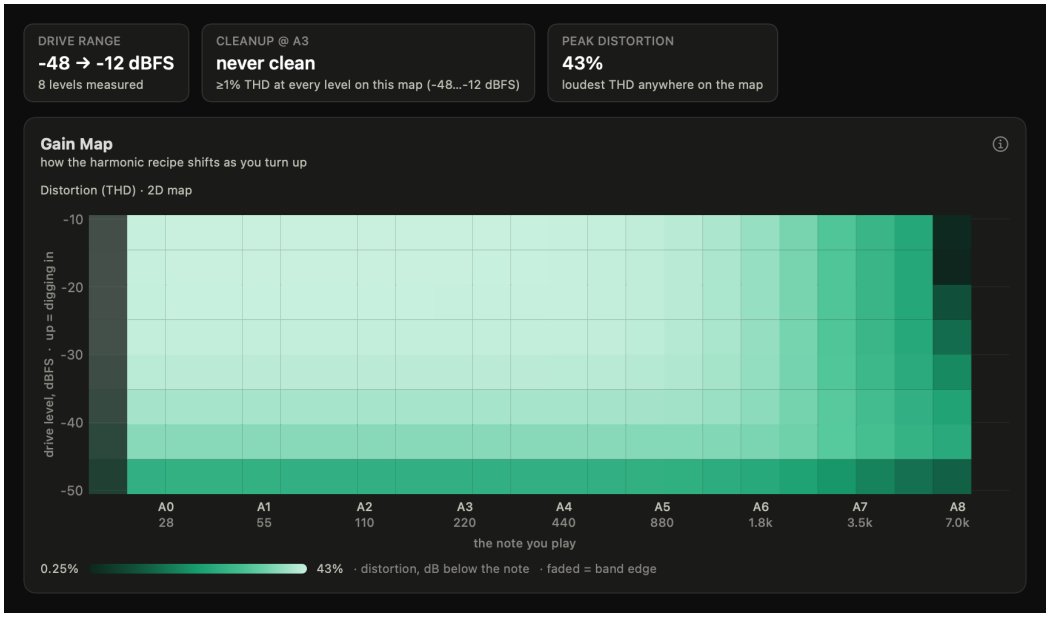


Figure 24: The Gain Map view as the app draws it, from the guide's simulated fixture: the THD map over note and level, the outermost columns drawn faded, quiet cells dimmed by their margin.

Every case is a known device through a known LOOP: a delay, one seeded noise stream referred to the interface's input, and a chain gain applied after the device and the noise, as an input-gain knob is; the calibration is the app's own sweep on the identity loop through the same gain and noise, analyzed by the shipped calibrator, whose *measured* latency cuts every level — the app's flow. Beside every cell the oracle prints its truth: the memoryless device's closed-form ladder at that level's drive (chapter one's quadrature, one pass for all orders), scaled by a pre-filter's magnitude at f_j, a post-filter's at kf_j and the chain gain, and the truth THD as (31) applied to that ladder; and the read-time half is the shipped predicates run on a summary built from the truth cells — the definition applied to the ladder, never a second derivation.

The instrument every truth comparison rests on is the leakage map (Figure 21): the pure phantom [1], whose truth is zero at every order above the first, read on the journey's geometry. Its higher-order cells are the analyzer's own leakage into each order's window per column — -70.13 dB into H_2's at the first unmasked column, -95.55 into H_3's, -123.5 into H_5's; a column up, -133.1 into H_3's:

k	26.2 Hz	34.33 Hz	44.99 Hz	58.94 Hz	173.7 Hz
2	-70.13	-103.6	-117.2	-127.2	-159.5
3	-95.55	-133.1	-152.3	-162.5	-201
4	-112.1	-145.6	-163.2	-176.4	-217.4
5	-123.5	-155.3	-169.3	-182.7	-221.4
6	-131.9	-162.7	-181.1	-194.2	-230.5
7	-138.3	-167.3	-188.3	-196.7	-240.2
8	-143.6	-171.3	-191.5	-196.7	-233.4
9	-148	-173.5	-197.6	-199.2	-235.6

A cell whose truth stands c dB above the leakage landing in its window can read up to $20\log_{10}(1+10^{-c/20})$ from the truth by the analyzer alone — 0.27 dB at $c = 30$, 0.0864 at 40 — and that allowance is subtracted before the bar is applied; what remains is the method's own error. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. The test asserts:

- Swift against Python: every header quantity and every content cell in decibels within 1e-06 (measured maximum 5.1e-08 on this run) on a capture with no filter in it, and within 0.0001 (measured 1.8e-06) through a biquad; every discrete reading — mask, presence, clear, corroborated, every tile's state, every column's cleanup state — identical; the stamp point for point; the residue per column per length wherever it is a reading (a column two hundred decibels down is dust on both sides, stated and not compared, as is a harmonic cell whose truth sits more than 140 dB under the fundamental);
- the shipped read against the ladder on every noiseless case, after the allowance: within 0.05 for the smooth devices (measured residue 0.012), within 0.05 for the twins (0.04), within 0.3 through a pre- or post-filter (0.15 — the truth treats each column's fundamental as a steady tone through the filter, and a sweep is not steady); the raw error tracks the bound: at 30 dB of clearance the worst cell reads 0.179 dB against the bound's 0.27, at 40 dB 0.0681 against 0.0864;
- the THD cells the same way, the allowance built from the residue at that column for the case's own sweep length: within 0.001 (measured 6.4e-05) and 0.2 (0.09);
- the cleanup crossing per unmasked column against the truth curve's crossing within 0.3 dB (measured 0.197, at the first unmasked column in every case, where the residue's H_2 leak sits inside the summed THD);

- the identity through noise at the three offered lengths: the residue at the first unmasked column reads the law's three figures from the Python's own analysis (-46.32 / -70.11 / -86.67 against -46.3 / -70.1 / -86.7, within 0.05), every cell's fundamental present, the peak tile a bound (0.133 % at the default length), the cleanup tile always clean, and no cell in the loudest row quoted clear;
- the #100 term: the same device and noise through a loop 12 dB louder moves H_1 by 12 dB, the stamp's chain gain to 12, and no composed margin by more than $\leq$ 1e-10 dB (bar 1e-06);
- a capture 100 samples late: the calibrator finds the latency and every harmonic cell moves by $\leq$ 1e-10 dB (bar 1e-09);
- Part B on the corpus grids: the Python's tiles against the kernel's, the peaks within 0.1 dB and the bounds to a relative 1e-09.

Part B's parity is established: on every read-time case — the identity through noise at the three sweep lengths, the device with no fundamental, the never-clean clipper, the loop 12 dB louder, the operative floor on a silent stamp, and the plugin window — the two implementations agree on every margin, every flag, every tile and every verdict, and on the corpus grids the reimplementation reproduces the kernel's own tiles to the digit.

The device cases, each with its content-cell count, its Swift-vs-Python figure, its worst raw error and where, its residue after the allowance, its worst raw error at the two clearance bars, and its THD residue:

case	cells	S−P (dB)	raw (dB)	at	residue (dB)	30 dB bar	40 dB bar	THD residue
a-tanh	426	5.1e-08	0.348	H5 (4,2)	0.012	0.0978	0.0681	4e-08
a-tanh-twin	479	3.3e-09	0.341	H5 (4,2)	0.0051	0.0981	0.068	4e-08
a-poly	456	3.5e-08	1.33	H3 (0,2)	0.00037	0.17	0.0534	6.4e-05
a-poly-twin	504	3.3e-09	1.52	H3 (0,2)	0.04	0.179	0.0545	1.7e-05
j-latency	426	5.1e-08	0.348	H5 (4,2)	0.012	0.0978	0.0681	4e-08
k-tanh-prehp	421	1.2e-06	0.215	H5 (5,2)	0.15	0.215	0.155	0.09
k-poly-postlp	456	1.8e-06	1.34	H3 (0,2)	0.00038	0.169	0.0544	7.1e-05

The read-time cases on the Swift side — the tiles the page would print — with the two implementations' agreement:

case	peak tile	value (%)	cleanup tile	level (dBFS)	absent	S−P (dB)
b-identity-2s	at floor (≤)	0.4834	always clean	–	0	3.3e-09
b-identity-5s	at floor (≤)	0.1332	always clean	–	0	3.3e-09
b-identity-10s	at floor (≤)	0.1245	always clean	–	0	3.3e-09
c-absent	no fundamental	–	none	–	1	3.3e-09
d-never-clean	peak	43.69	never clean	–	0	3.3e-09
e-tanh-gain0	peak	1.98	cleans up	-15.74	0	5.1e-08
e-tanh-gain12	peak	1.98	cleans up	-15.74	0	5.1e-08
f-operative	peak	1.98	cleans up	-15.74	0	5.1e-08
g-plugin-hardclip	peak	41.57	cleans up	-25.05	0	3.3e-09

The worked example, $y = \tanh(2x)$ noiseless on the hardware lattice: 426 content cells compared, the worst raw error 0.348 dB at H5 (4,2) and 0.012 dB after the allowance, 0.0978 dB at the 30 dB bar and 0.0681 at 40, a THD residue of 4e-08 dB; its peak tile reads 1.979 % against the truth grid's 1.979 %, and its cleanup tile -15.74 dBFS at 227.6 Hz against the truth curve's -15.74 (a difference of 0.044 dB). The device with no fundamental — a phantom whose fundamental sits 80 dB under its second harmonic — reads a majority verdict of absent (0.79 of the cells) at a median depth of 80.03 dB and a tile with no percentage; the cells beside the low edge read present, because the fundamental's window there holds the second harmonic's own leakage. The

operative-floor case: the render-path stamp reads -101 dB at its loudest point, the stated floor enters at -76.99 dBFS, and the peak reads 1.98 %. The plugin window with the hard clipper: a peak of 41.57 % and a crossing at -25.05 dBFS. The never-clean clipper: 43.69 % at every level of its tile column.

Part B on the corpus — the kernel test's own transcribed grids and stamps, the read-time half's one stored-input exception (a stored record holds results, not a capture, but for the read-time half the stored grid *is* the input, and these grids are the store's own figures pinned to the digit by the kernel's test):

record	tile	Python (%)	kernel (%)
59	at floor (≤)	0.964566	0.964566
428	at floor (≤)	0.778677	0.778677
65	peak	11.2116	11.2116
446	peak	6.21905	6.21905
404	at floor (≤)	0.130628	0.130628
410	at floor (≤)	0.106096	0.106096
416	at floor (≤)	0.119273	0.119273
422	at floor (≤)	0.130136	0.130136
464	at floor (≤)	0.105552	0.105552

6.8 Limits and failure modes

- **The sweep length is not stored, and the floor pays for it.** The residue is taken as the worst over the offered lengths, so a record captured at the default length carries the short sweep's residue at the first unmasked column, -46.32 dB against its own -70.11: that column's margin is about 24 dB pessimistic — a bound in the honest direction, and a stored length would sharpen it (a stored-format decision, not a read-time one).

- **The fundamental grid is a level, not a normalizer.** The cells are uncompensated, so $h[1][i][j]$ is the fundamental as the device delivered it through the loop, gain and all. On a device with a series resistance ahead of its clipper it is also a load gauge: the same reference position measured on two interfaces with different input impedances delivers a different fundamental, and every ratio to it is taken against that delivered level. The map does not separate the device's gain from the loop's, and it does not claim to.

- **The analyzer's separation is geometry-dependent, and the journey's geometry is the coarser one.** The default sweep is half the fingerprint sweep's length and the response DFT half its length, so the leakage of the loudest order into each window at the low edge is worse than chapter one's measured bound; the leakage map states it per column and order, and a cell whose content sits inside it is the analyzer's, not the device's — which is what the residue and the margins say in a real capture, and why a device with no fundamental reads present beside the low edge.

- **Aliasing in a simulated device.** A device applied per sample folds its content above Nyquist onto the reads; the smooth devices' residues above are inside the allowance, the phantom twins are the alias-free control, and a kinked device (a hard clipper at playing levels) would not be the oracle for a cell near the top of a band.

- **A raw maximum is a peak where no floor exists.** Without a calibration or an operative floor every unmasked cell counts as clear and corroborated, and the tile states the map's loudest cell as a peak — Simulation Mode and every record with no calibration are in that state; reported, not changed.

- **The per-harmonic path reads the stamp verbatim.** The §3 distance and the fundamental-presence margin read the calibration stamp point by point; the single-read exposure the median read removes from this map exists there and is reported, not changed — moving every distance figure in the library is its own measured change.
- **The audition models invent.** The per-level fits are consumed by the listening panels only, and a render between two levels blends a clamped quieter fit into the louder one; nothing measured reads a fit, and the invention's size is a later chapter's measurement.
- **The rows are separate captures.** Each level is its own sweep seconds apart; a circuit whose operating point moves between them shows it as row-to-row wobble, which the map draws and does not name.

7 Chord IMD

7.1 Significance and limitations

A two-tone measurement plays two sines at a musical interval through the device at once. A nonlinearity does not treat them separately: it multiplies them, and the output carries every intermodulation product $mf_1 + nf_2$ for integer m, n, whose product order $|m| + |n|$ is the lowest polynomial power that can produce it. The result is the level of every product up to the enumeration ceiling K (the parameter table gives it, and the count of recipes it admits), each read at one bin of one long spectrum and quoted in dBc against the louder played tone; the two tone amplitudes as the device delivered them; three capture-time percentage sums (two borrowed from laboratory two-tone standards and one total, kept for the record as the frozen tile definitions of earlier builds); and the raw spectrum beside each played tone, stored for the coherence check. At read time every product gets one product reading, the capture states its own capture floor, the products are sorted by harmonic role on the grid of the requested interval, and the drive level is stated from the stored stimulus.

What it cannot tell you, stated once: (a) anything about other drive levels — one capture is one operating point, and the products of a clipper scale with drive faster than anything else about it; (b) whether a resolved product is the device's or the rig's — an interface's converters intermodulate too, and only a control capture with nothing in the loop separates the two, which the floor cannot do and does not claim to; (c) which stage moved the notes when the coherence check fails — a verdict without an attribution, by design (§7.4); (d) anything below the capture's own floor, and, for the three percentage sums, anything at all about a floor: they are capture-time sums over the products the gate passed and carry no bound (§7.8); (e) anything about products above order K — the enumeration is a small-signal census, and a hard clipper at playing levels has energy at orders it does not walk; and (f) anything about what the distortion sounds like.

7.2 The stimulus

In this chapter $f_1 < f_2$ are the two played tones (not the sweep band edges of §3), N is the analysis window length in samples, $\Delta f = f_s/N$ the bin width, k_1, k_2 the bins the two tones sit on, K the enumeration ceiling, and (m, n) a product's recipe. The stimulus is two equal sines at a linear combined peak A:

$$x[n] = \tfrac{A}{2}\sin(2\pi f_1 n/f_s) + \tfrac{A}{2}\sin(2\pi f_2 n/f_s), \qquad n = 0 \ldots \lfloor T f_s + \tfrac{1}{2} \rfloor - 1, \tag{35}$$

scaled over its first and last $\lfloor t_{\text{fade}} f_s \rfloor$ samples by the raised cosine $\frac{1}{2}(1 - \cos(\pi i/N_{\text{fade}}))$. The standard peak is the fingerprint sweep's drive, so the products and the other measurement kinds describe one operating point; the worked example below uses the loudest drive the hardware sheet offers, where the reference clippers clip. The notes are named in twelve-tone equal temperament with A4 = 440 Hz (the parameter table gives the default pair), and the measure sheet offers one probe: the interval. Two laboratory probes — an SMPTE-style pair at a fixed ratio and a CCIF twin-tone — existed in the code and were retired before 1.0, because a standard is never moved onto the lattice described next and neither could be given one; the two percentage keys that carry their names are sums over recipe sets and are unrelated to the retired probes.

Coherent sampling, the premise — stated once. Both tones are moved onto exact bin centers of the N-point window before anything plays, so each completes a whole number of cycles inside the segment the analyzer reads, and the segment is windowed *rectangularly*: a tone that fits the window exactly leaks into no other bin, so there is nothing to taper and no taper is used, and the floor beneath a record is noise, not a window's skirt. coherent sampling is inherited by every product: with the tones on bins k_1, k_2 the product (m, n) lands on bin $mk_1 + nk_2$ exactly, because

bin numbers add as frequencies do. That set of bins is the analysis lattice. Every step below is a consequence of this one premise, and the price of it is the coherence check: a rectangular window's sidelobes fall as $1/(\pi d)$ at d bins — excellent, with no margin — so a tone that drifts during the capture spills its own skirt across the lattice.

Two numbers describe the lattice. Write $k_1 = ga$, $k_2 = gb$ with $\gcd(a, b) = 1$: the reduced tone ratio $a : b$ and the lattice spacing $g = \gcd(k_1, k_2)$. They answer two different questions. Two recipes share a bin when $(m - m')\,a = -(n - n')\,b$, which with a, b coprime forces $m - m' = -jb$, $n - n' = ja$ for some integer $j \neq 0$, so $|m - m'| + |n - n'| = |j|(a + b) \leq 2K$; hence

$$a + b \;\geq\; 2K + 1 \tag{36}$$

admits no collision up to order K (the bound is tight: at $K = 7$ the ratio $1 : 13$ collides and every ratio at 15 is clean). And the integer combinations of k_1 and k_2 are exactly the multiples of g (Bézout), so two distinct lattice points differ by at least g bins, and some pair by exactly g: no product can sit nearer than g bins to another product or to either tone, at any order. A lattice can pass (36) and still be crowded — the probe every record before the current length was captured on shared no bins at a spacing of two — which is why the snap buys spacing and not only collision-freedom.

The lattice snap. The lattice snap is decided in two steps, so that the interval a record asks is one interval at every rate and length. First the ratio is chosen *once*, against every standard grid at the same time (the four supported rates, each at the length it is meant to analyze at): over $a = 1 \ldots \lfloor f_1\rho/(\beta\,\Delta f_{\max})\rfloor$ with $\rho = 2^{c/1200}$ the cents budget c as a ratio and β the resolvable-spacing bar, and for each a over $b = \max(a+1, \lceil ar/\rho^2\rceil) \ldots \lfloor ar\rho^2\rfloor$ with $r = f_2/f_1$, a coprime pair satisfying (36) is a candidate; it survives when on *every* standard grid some multiple $g \geq \beta$ places both $ga\,\Delta f$ and $gb\,\Delta f$ inside $\pm c$ of the requested tones (the placement on one grid is the in-budget g with the smallest worst-case pitch error, ties to the smaller g). The survivors are ranked by the interval nearest the requested one, then the wider minimum spacing, then the smaller $a + b$, then (a, b). The default fifth lands on the ratio the parameter table quotes, a few cents flat of equal temperament and the same at every rate; the snap table in §7.5 is that search run by the shipped code on every grid. Second, on the grid at hand, the standard ratio is scaled: the in-budget multiple with the smallest pitch error at *any* spacing — keeping the interval outranks clearing the bar, so a grid too coarse to resolve it (Simulation Mode's interactive preview) still asks the same musical question, in a degraded regime the read-time predicates carry. Only when no multiple of the ratio fits the budget does the older search run: every bin pair (k_1, k_2) inside the budget (the half-widths searched are derived from the budget and the bin width, never fixed), ranked by spacing, then worst-case cents, then index; the best pair clearing both (36) and $g \geq \beta$ is emitted as *resolvable*, else the best clearing (36) alone as *degraded* — which at the old length is the corpus probe, field for field — else the plan *refuses* before anything plays, quoting the widest lattice the budget reached with its reduced ratio, spacing and sum against the bound. The plain snap is never returned as a fallback: an exactly degenerate probe drops colliding products through the enumeration's tie-break and reports a suspiciously clean record.

The window is a duration, and the stimulus is sized to hold it. The standard analysis length is stated at the standard rate as a sample count, but what it fixes is a *duration*: each rate takes the power of two nearest $N_{\text{std}}\,f_s/f_{s,\text{std}}$, capped at the plan's ceiling, so the four rates form two grid classes with the same bin width in pairs. The stimulus duration is then derived from the length it must hold,

$$T = t_{\text{fade}} + t_{\text{settle}} + N/f_s + t_{\text{fade}} + t_{\text{settle}}, \tag{37}$$

never set by hand: a hand-set duration is how an earlier fixed window came to overrun the fade-out at two of the four rates (Figure 26). The analysis length a stored stimulus is read at is derived by the same rule in reverse — from the ceiling, halve while the window would overrun the steady state — so the stimulus is snapped to exactly the length it is analyzed at, and a shorter snap grid is a subset of a longer window's bins, never the reverse.

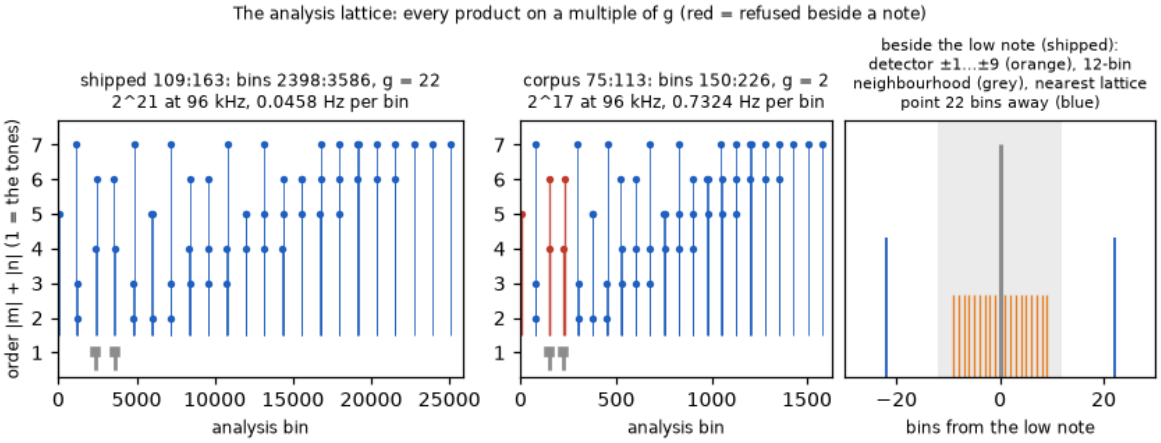


Figure 25: The analysis lattice: every enumerated recipe as a stem at its bin and order on the shipped fifth (left) and on the corpus probe every earlier record was captured on (middle), the two tones as squares at order one. On the corpus probe the four even-order recipes two bins from the played notes and the one two bins from DC draw red: refused on position (§7.4). Right: the bins beside the low note on the shipped lattice — the detector's offsets, the played-tone neighborhood, and the nearest lattice point a spacing away.

7.3 The pipeline

Given the captured output $y[n]$ and the integer loop latency ℓ (the app measures and removes it per capture; the oracle renders at zero latency and, when asked, analyzes with a deliberate offset):

1. **The window.** The segment starts at $s = \ell + \lfloor t_{\text{fade}} f_s \rfloor + \lfloor t_{\text{settle}} f_s \rfloor$ and runs for exactly N samples, rectangularly windowed and zero-padded if the capture ends first; the steady state ends at $e = \ell + \lfloor T f_s + \frac{1}{2} \rfloor - \lfloor t_{\text{fade}} f_s \rfloor$, and the *overrun* $s + N - e$ is positive when the segment would swallow the fade-out. Latency shifts s and e alike, so a window that fits at zero latency fits at any, and because the tones are bin-centered a shift is a phase: no magnitude moves (§7.7 measures it). The analyzer never shortens a window it is handed — the tones are centered on the length the stimulus was snapped to, and a shorter rectangular window would leak far worse than the fade it avoided.

2. **The spectrum.** With c the number of captured samples inside the segment,

$$M[b] = \frac{2}{\min(N, c)} \left| \mathrm{DFT}_N\{y[s+i]\}[b] \right|, \qquad b = 0 \ldots N/2, \tag{38}$$

   so a bin-centered sine of amplitude a reads a.

3. **The enumeration and the lattice.** With $\mathrm{bin}(f) = \mathrm{round}(f/\Delta f)$, the set $S = \{\mathrm{bin}(f_1), \mathrm{bin}(f_2), 0\}$ is seeded, and the walk runs order $= 2 \ldots K$, $m = -\text{order} \ldots \text{order}$, $n = +(\text{order} - |m|)$ then $-(\text{order} - |m|)$: a recipe with $\Delta f < m f_1 + n f_2 < 0.95\, f_s/2$ (both edges literals) whose bin is not yet in S is admitted and its bin joins S. The lowest order

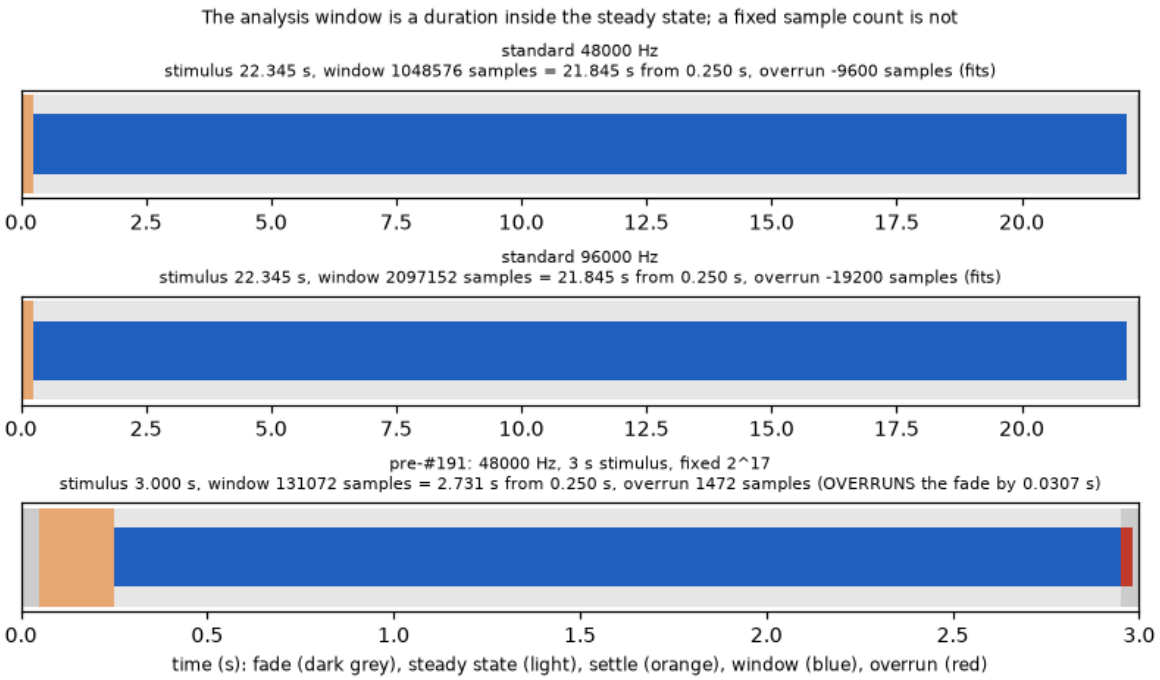


Figure 26: The analysis window inside the stimulus at the two 96 kHz-class rates, and the historical fixed-length window on a hand-set stimulus at 48 kHz, which overran the fade-out: periodicity broke and the window's own skirt flooded the low bins. The fade-contaminated separation bound in the parameter table exists for captures made that way.

therefore keeps a shared bin — on a degenerate lattice a surviving label names the enumeration's tie-break, not a fact — and a clean lattice admits $\sum_{k=2}^{K} 2k$ recipes. The lattice is collected *before* any floor is read, so a floor's neighborhood can leave out every bin the enumeration claims.

4. **The reads.** Each product's amplitude is $M[\mathrm{bin}(mf_1 + nf_2)]$, the one bin, exactly; the tone amplitudes are $M[k_1]$ and $M[k_2]$. An older read took the largest of the bin and its two neighbors "to tolerate drift" there was none of: it left a strong product untouched and lifted every read at the noise by a measured mean of about four decibels, and the separation bound in the parameter table is deliberately left where it was measured under that read, because a stored record may carry it.

5. **The local floor and the gate.** Each product is judged against the median magnitude of its own neighborhood: the bins from the inner offset to the outer offset on both sides, inside the spectrum, with every bin the lattice claims removed by name — every product bin, both tone bins, DC, and the coherence detector's bins beside each tone. The inner offset is a bin quantity (the product's own skirt under lost coherence); the outer offset is a *frequency* span converted to bins at the record's bin width and never fewer than a minimum count (a noise density is local in hertz, and a median wants samples). The median is the upper middle of the sorted values, and zero when the neighborhood is empty. The presence gate is

$$\text{present at capture} \;\Leftrightarrow\; a_{m,n} > 4\,\mathrm{floor}_{m,n}, \tag{39}$$

a literal factor of four (the table states it in decibels), decided once and *frozen* in the payload as the product's flag. That is the local floor; the capture floor of §7.4 is a read-time summary that can never change it.

6. **The percentage sums.** Over the products inside the audible band B (the band is stored

with the record),

$$\begin{aligned} \mathrm{pct}(P,R) &= 100\sqrt{P/R} \quad (R>0),\\ \mathrm{SMPTE} &= \mathrm{pct}\Big(\sum_{|n|=1,\ 0<|m|\le 4} a_{m,n}^2,\ a_{f_2}^2\Big),\\ \mathrm{CCIF} &= \mathrm{pct}\Big(\sum_{(m,n)\in\{(-1,1),(2,-1),(-1,2)\}} a_{m,n}^2,\ a_{f_1}^2+a_{f_2}^2\Big), \end{aligned} \tag{40}$$

and the total is the same over every in-band product the gate passed. These are the capture-time scalars; the read-time derivation of §7.4 re-runs them with the lattice's refusals left out.

7. **The coherence detector.** For each tone bin the raw magnitude at every detector offset on both sides, inside the spectrum, and the local floor beside the tone (the same neighborhood and exclusion about the tone bin) are stored with the window length: the coherence detector. Nothing else keeps those bins, and the read-time check needs no derived length to place them.

8. **Storage.** The products are stored sorted by descending amplitude with their recipe, frequency, amplitude, order and gate flag; beside them the stimulus as delivered, the two tone amplitudes, the three sums, the band and the detector. Nothing read-time is persisted.

7.4 Read-time predicates on the result

Everything below is computed from the frozen payload when a record is read, so a corrected predicate corrects every stored record. A consumer that knows the analysis length passes it; one that does not passes nothing, and nothing is placed on the lattice for it — an absent argument is a smaller claim, stated rather than defaulted.

Levels and the capture floor. A product's level is $L = 20\log_{10}\left(a/\max(a_{f_1}, a_{f_2})\right)$ dBc, undefined when either amplitude is zero. The *noise reads* are the levels of every product the gate did *not* pass and that the lattice does not refuse — not inside a reference bin's neighborhood, not under the band's low edge, not beside a louder product (each defined below) — sorted ascending; bins under the loudest product's skirt deliberately stay in, because on a record like a rectifier's every empty bin is under it and a vanished floor would read as "every product resolved". With at least the minimum count of reads the capture floor is the upper median, with the count and the spread of the loudest read above the median; fewer reads is honestly no floor. The two floors run one way: the gate is capture-time against the raw spectrum, the capture floor is read-time over bins already flagged, so excluding a refused bin changes which flagged bins are averaged and never which are flagged.

Positions on the lattice. *Proximity*: a product's bin against the three reference bins — $\mathrm{bin}(f_1)$, $\mathrm{bin}(f_2)$ and 0 — the nearest wins (ties in that order), and the product is inside the neighborhood when the absolute offset is at most the played-tone neighborhood width (a bin quantity: sidelobes fall as $1/(\pi d)$ in bins). *Band*: the stored frequency against the stored band's *low* edge only; a product above the band's top is inaudible but perfectly attributable and reads like any other. *Louder neighbor*: among the products that passed the gate and are strictly louder, those within the same width of this product's bin (offset nonzero) — the nearest, ties to the louder; a neighbor that is itself noise casts no skirt and refuses nothing. *Spacing*: $g = \gcd(\mathrm{bin}(f_1), \mathrm{bin}(f_2))$ on the analysis length. *Loudest skirt*: a record-level condition, present only when the loudest *product* (the maximum of L over every stored product, never a quantity that folds the tone in)

stands above 0 dBc and the coherence reading is absent or does not interpret for a device; it carries g and the skirt depth $20\log_{10}(\pi g)$, computed never typed. The *analyzer window* bound is the fade-contaminated separation of the parameter table, asserted only when the window overruns the fade (§7.3, step 1) and otherwise nil — a record captured at the old fixed length at one of the two rates where it overran gets that bound; every capture the app makes now gets none.

The reading. One product reading per product, decided in this order and stopping at the first that applies:

1. inside a reference bin's neighborhood (a played tone, or DC) — *cannot resolve, beside a played tone*, naming the tone and the signed offset: a bin inside a note's own skirt holds the note, and a bin beside DC holds the capture's own offset and drift; decided on position before any level is consulted, because the gate samples the floor from offsets that lie mostly outside the skirt the bin is sitting in, and would pass such a bin by its own rule at tens of decibels of margin;
2. under the band's low edge — *out of band*: a frequency fact, not a limit, so its sentence never merges with a lattice's or the instrument's;
3. no level (a zero amplitude or reference) — *unmeasurable*;
4. a louder neighbor inside the neighborhood — *cannot resolve, beside a louder product*, naming it: at that separation the quieter bin holds the louder product's skirt; the louder of the pair is readable, the quieter is not;
5. depth below the loudest component at or beyond the analyzer's bound, when one is asserted — *cannot resolve, analyzer window*;
6. deeper below the loudest product than its skirt reaches, on a record where the loudest-skirt condition holds — *cannot resolve, under the loudest skirt*: unread, never present and never absent;
7. the gate passed — *present*, with the margin over the capture floor when one exists;
8. no capture floor — *cannot resolve, floor not observed*, carrying the number of noise reads;
9. otherwise — *absent*, with the bin's own level as the bound: nothing here above it, a bound and never a value.

A refusal is never silently absent, and "absent" is never a value. The verdict counts present, absent, unresolved (any *cannot resolve*) and out-of-band readings with the noise-read count; a nil floor is stated by its reason — *every product resolved* only when nothing was refused, else *too few empty bins* with the counts — in the one floor sentence every surface prints.

Shared bins, and the lattice a record was delivered on. Every capture path snaps both tones, so distinct lattice points differ by at least one bin width while collisions agree to the ulp; a stored product's bin is *contested* when another enumerated recipe (the default ceiling raised to the record's own highest stored order, since the ceiling is not stored) lands at the same frequency to a relative tolerance, and the contestants are listed ascending in order. The delivered lattice — the reduced ratio of the two stored tones and the grid f_1/a — is recovered from the stored signal by the smallest denominator that reproduces the pair to the same tolerance, which no snapped record can fail.

The derived percentages and the partner check. The three sums of (40) are re-run over the stored products with the recipes inside a reference bin's neighborhood left out, the band as stored (every product on a legacy record with no band), and the excluded list names only what the band would have counted — the recipe beside DC sits under the band and was never in any sum. The partner check is the second-order cross-check: a quadratic term produces $f_1 + f_2$ with $f_2 - f_1$ at equal level at any tone balance, and $2f_1$ with $2f_2$ at twice the measured *output* tone gap; only *excess* over the partner beyond the tolerance flags, a contested operand refuses, and a flag says the bin holds something beyond a clean second-order product — the energy is real, the label and rank are not to be trusted.

The coherence check. The detector's bins are read as levels in dBc, each pair (the sidebands about one tone) against the local floor beside that tone in the same unit. On a record with no detector — every capture before the detector existed — the same reading is built from the stored products inside a played tone's neighborhood (never DC's) against the capture floor at the derived length: the legacy path, which the detector outranks whenever present. The lattice facts: g on the detector's own length, the widest stored offset, and the *clearance* $g - d_{\max}$; the lattice is *clear* when the clearance is at least the analyzer's inner offset. The detector offsets are $\pm 1 \ldots \pm d_{\max}$ with $d_{\max}$ *derived* as the resolvable-spacing bar minus one minus the inner offset — the widest offset that stays the inner-edge clearance from the nearest lattice point on the narrowest lattice the snap calls resolvable — and every offset is read, odd and even: a wander at an even number of cycles per window puts its sidebands on even offsets only, and an odd-only detector read the floor on the very fixture whose stored two-bin products read tens of decibels above it.

Clearance is necessary and not sufficient, because distance is a proxy for level. Both a tone and a product spill a $1/(\pi d)$ skirt under lost coherence, so a bin at offset d beside a tone reads the tone only while the loudest product, assumed at the nearest lattice point $g - d$ bins the other way, sits under the tone's skirt there by the stated margin μ:

$$L_p \;\le\; 20 \log_{10} \frac{g - d}{d} - \mu \;=\; \mathrm{ceiling}(d, g), \tag{41}$$

$-\infty$ when a lattice point sits at or inside the offset. The *readable* offsets are those whose ceiling the loudest product clears (every offset on the legacy path or an uncleared lattice); the ceiling falls with d, so the set is always $\pm 1 \ldots \pm k$ and the nearest bins survive last. The figure is the loudest sideband over the readable offsets of each pair, falling back to every sideband only when the level rule withheld every offset (and then the reading does not interpret); the excess is that figure over its pair's floor, the worst over the two pairs, and the bar is a floor-relative excess (the table's value) — an absolute bar was right at one length only, since spilled carrier energy does not scale with bin width while noise drops with it. Beside the excess the line prints what noise alone would put the worst of the n read bins over the floor: a noise bin's power is exponential, the floor a median, and the expected worst of n over the median is

$$10 \log_{10}(H_n / \ln 2), \qquad H_n = \textstyle\sum_{i=1}^{n} 1/i, \tag{42}$$

the exact form, chosen over the $\sqrt{2 \ln n}$ asymptote which under-reads by nearly two decibels at the legacy path's four bins — a reference, never a bar. The sideband pair statistic is the mean of the two innermost sidebands of each pair and the difference between the two tones' pairs, printed beside $20 \log_{10}(f_2/f_1)$: a shared fractional frequency wander separates the pairs by that figure, while amplitude modulation or a frequency-independent phase modulation leaves them equal — a quantity the sentence states, and no verdict is drawn from it (equality is never read as amplitude modulation, a biconditional the shape has not earned).

The reading *interprets* when a floor exists and either the lattice is clear with at least one readable offset — a statement about the played tones whatever was in the loop — or the lattice is not clear

and the subject is a loop-only capture, which held nothing that could put a product beside a note. The subject (loop only, or a device in the loop) is stated by the caller, which knows the record's identity; it is context on a cleared lattice and the deciding fact only on an uncleared one. The verdict is about the *notes*: when the reading interprets and exceeds the bar, the played notes did not stay stationary across the window, from some source in the chain the reading cannot name — a pedal with a dying battery moves a note over twenty seconds as surely as a rig. Nothing is refused, discarded or re-run; the capture is saved as measured and the sentence rides it. Where the level rule withheld every offset the sentence names the signal that would have been read instead and draws no verdict; where the lattice is not clear and a device was in the loop the number is shown and labeled not attributable.

What the verdict means for each product. A note that did not stay stationary spills a skirt, and a rectangular window's skirt falls as $1/(\pi m)$ in bins, so what reaches a product bin twenty-two bins from a note is a fraction of what the detector read beside it: on the published day-B set the sidebands sit near -70 dBc within nine bins of a note and the loud products at $-5 \ldots -57$ dBc, which the smear cannot touch to a fraction of a decibel, while the floor near -107 dBc and the products near it are genuinely unreadable. The reading therefore does not declare every product untrustworthy; it predicts the smear at each product's bin and reads the products that stand clear of it. Two measured facts decide the model (the validated wander generator through five devices, five wander sizes and five shapes, with a torus-quadrature truth beside every product): the far skirt follows the *outer* detector reads, not the inner ones — twenty-two bins from a tone the smear read -121 dBc where a $1/m$ envelope anchored on the ± 1 read predicted -102, because near a modulation sideband the Dirichlet kernel is not yet in its $1/x$ regime — and a loud product casts a skirt of its own, decisive at the 17–29 dB level: a product $mf_1 + nf_2$ of a modulated pair carries the modulation with an index scaled by its frequency (a frequency wander) or its order (an amplitude wander before the clipper). So each tone's skirt at distance D is

$$S_{\text{tone}}(D) \;=\; \max_{d \ge d_a,\; d < D} \big[\, L_d + 20 \log_{10}(d/D) \,\big], \tag{43}$$

over its stored reads L_d at offsets from the anchor d_a (the table's value) outward, both sides; each product that stands above its local floor contributes its own level plus the tones' relative skirt at D plus $20 \log_{10} \max(\text{order},\, f_p / \max(f_1, f_2))$ — the larger of the two mechanisms' scalings, so no mechanism is inferred; the contributions are summed in power and the dominant source named. A product is *read* where its bin stands at least the smear margin (the table's value) above the prediction, and *unread under the smear* otherwise — not present, not absent, its tile a lower bound, and a quiet bin under it left out of the floor. The margin is derived from the same runs: the smallest 5 dB band at which the measured worst level error on a read product stayed under 2 dB (1.9 dB at 10–15 dB of margin, 2.8 at 5–10, up to 29 below zero), with a coherent worst case of $+2.4/-3.3$ dB at the margin. The sentence states the counts — how many products are read clear of the smear and how many are unread under it. A record whose check holds, a reading the level rule withheld entirely, an uncleared lattice and the legacy stored-products path carry no field and read as before; the legacy path and an uncleared lattice keep the first monitor's blanket sentence, since a loop-only capture has no product to read.

The harmonic roles. The tiles partition the products by what each does to the chord that was *requested* — classified at the requested interval, measured at the lattice. The request is not on the record, so the nominal notes are the nearest equal-tempered pitches to the delivered tones, unambiguous inside the snap's cents budget; this is the one place the delivered frequencies are consulted, and only to name the notes. The interval's just ratio $a : b$ is the simplest five-limit ratio (every prime factor 2, 3 or 5) within the pitch tolerance of the equal-tempered interval, found by search — smallest $a + b$, ties to the nearer cents — five-limit because a seven-limit ratio

can sit nearer by cents and put the grid's fundamental somewhere the music never goes. The grid's fundamental is f_1/a (55 Hz for the default fifth at $2:3$; 27.5 Hz for the major third at $4:5$), and a recipe lands on grid index

$$k = ma + nb, \tag{44}$$

exact in integers. Its pitch class relative to the lower note is read from the odd parts (factors of two stripped — octave reduction, exact in integers): cents $= 1200 \log_2(\text{odd}(k)/\text{odd}(a)) \bmod 1200$, the nearest semitone, the deviation from it, and the class modulo twelve; a point is a chord tone when $\text{odd}(k)$ equals $\text{odd}(a)$ or $\text{odd}(b)$, and between the frets when its deviation exceeds the tolerance. The roles: a chord tone at $k \geq a$ is *Thickening*, below *a Growl (sub-octave)*; otherwise a point between the frets is *Sourness (clash)*, a consonant class (the parameter table lists them; the ninth is the one deliberate departure from the classical list) is *Added colour*, and any other class is *Sourness (clash)*; $k \leq 0$ has no role. The tolerance is derived from the grid: the widest gap between a five-limit just interval and its equal-tempered class is under it and the nearest higher-limit point a two-tone lattice produces is beyond it, an empty band with the bar inside. The role table in §7.5 is the shipped classifier run on the default fifth's grid. A role's figure is the power sum, in dBc re the louder played note, of its present, in-band products counted *once per analysis bin* (loudest first, so a duplicated bin keeps its louder read); an absent product contributes nothing and is bounded by the page's floor; a product the instrument could not read was never looked at, so a role that holds one is a *lower bound* on its energy, stated beside the figure; a product under the band, or at grid index at most zero, is placed nowhere and marks nothing. *All products* is the same sum over every role. The record also states what it enumerated — the recipe count and the orders spanned — so a tone absent from the list can be told from one that was measured and silent.

Drive level and summaries. The drive line is read from the frozen stimulus: the combined peak $a_1 + a_2$ and each tone's own amplitude in dBFS ($20 \log_{10} \max(a, 10^{-12})$), volts only when the record's own embedded calibration carries the factor (the per-tone volts rounded to the label's digits, the combined figure derived from the rounded per-tone figures so the printed arithmetic holds), and a note that the two tones are equal — every order relation above assumes it. The summaries consume the present products at the derived length and one ratio of the cross products to the harmonic ones; a record that resolved nothing says so through the same verdict.

7.5 Parameters

Every value below is read from the manifest, which read it from the shipped code; the gate's factor, the enumeration's frequency edges, the pitch reference, the just-ratio search's numerator span, the contest tolerance and the spectrum's scale are literals the manifest carries as rule strings.

Parameter	Value
Plan	
default interval	A2 + E3
default combined peak A (re full scale)	0.05 (-26.02 dBFS; each tone $A/2$)
standard rate / supported rates	96 000 Hz / 44 100, 48 000, 88 200, 96 000 Hz
standard analysis length (the ceiling) N	2 097 152 (2^{21})
minimum analysis length	4 096
fade-in and fade-out t_{fade}	0.05 s
settle t_{settle}	0.2 s
The lattice search	
tone cents budget	±15 cents
resolvable spacing bar (g ≥)	13 bins

Parameter	Value
enumeration ceiling (max order K)	7
minimum reduced sum $a + b$ at K	15
delivered-lattice denominator limit	8 192
standard ratio $a : b$ (the fifth)	109 : 163 = 696.653 cents (requested 700; 26 ratios survived; runner-up 105 : 157 at 696.45)
The analyzer	
enumerated recipes (orders 2–K)	54
local floor: inner offset / span / minimum offsets	3 bins / ±8.789 Hz / 12 per side
local floor: outer offset at 2^{21} (96 kHz) / 2^{17}	192 / 12 bins
presence gate	amplitude > 4× the local floor (12.04 dB, a literal)
enumeration ceiling in frequency	0.95 $f_s/2$ (a literal)
The read-time predicate	
local floor separation	12.04 dB
minimum floor samples	3
fade-contaminated separation bound	32.5 dB
played-tone / DC neighbourhood	12 bins
skirt depth $20 \log_{10}(\pi g)$ at $g = 22$ / 24	36.79 / 37.55 dB
partner-check tolerance	6 dB
The coherence check	
minimum clearance (the inner edge)	3 bins
detector offsets	$\pm 1 \ldots \pm 9$ bins (derived: $13 - 1 - 3$)
bar over the local floor	30 dB
product skirt margin	10 dB
per-offset ceilings at $g = 22$ (dBc)	16.4, 10, 6.03, 3.06, 0.63, -1.48, -3.38, -5.14, -6.81
expected noise excess, 4 / 36 bins	4.78 / 7.8 dB
smear margin (a product is read this far above the predicted smear)	10 dB
smear anchor (the innermost detector offset the far skirt is extended from)	6 bins
The harmonic roles	
pitch tolerance	20 cents
added-colour semitone classes	2, 3, 4, 5, 7, 8, 9
just-ratio search limit	32 (denominator; numerator to 16×)
Drive level and pitch	
volts label digits	3
pitch reference	A4 = 440 Hz, 12-TET (a literal)
audible band	20 Hz – 20 000 Hz
Oracle (`analysisdump imd-synth`)	
torus quadrature points per axis	4 096

The lattice the shipped search chooses on every grid — the standard ratio placed on each, and the fallback on the two grids too coarse to place it — and the major third at the standard grid; candidates are the resolvable and the non-degenerate counts the search saw:

probe	f_s (Hz)	N	tier	$k_1 : k_2$	g	$a : b$	cents	candidates
fifth 44100 Hz	44 100	2^{20}	standard resolvable	2 616 : 3 912	24	109 : 163	0.334 / -3.01	121 / 3 061
fifth 48000 Hz	48 000	2^{20}	standard resolvable	2 398 : 3 586	22	109 : 163	-3.6 / -6.94	109 / 2 523
fifth 88200 Hz	88 200	2^{21}	standard resolvable	2 616 : 3 912	24	109 : 163	0.334 / -3.01	121 / 3 061
fifth 96000 Hz	96 000	2^{21}	standard resolvable	2 398 : 3 586	22	109 : 163	-3.6 / -6.94	109 / 2 523
fifth corpus 96000 Hz / 2^17	96 000	2^{17}	degraded	150 : 226	2	75 : 113	-2.15 / 7.48	0 / 8
fifth Simulation Mode interactive	48 000	2^{15}	degraded	75 : 113	1	75 : 113	-2.15 / 7.48	0 / 2
major third A2 + C#3 96000 Hz	96 000	2^{21}	standard resolvable	2 400 : 3 024	48	50 : 63	-2.15 / -2.04	99 / 2 127

The window on every grid, from the shipped geometry on the snapped default probe (sample indices; a negative overrun is the margin before the fade-out):

grid	f_s (Hz)	N	bin (Hz)	stimulus (s)	start	steady end	overrun
standard 44100 Hz	44 100	2^{20}	0.042057	24.2772	11 025	1 068 421	-8 820
standard 48000 Hz	48 000	2^{20}	0.045776	22.3453	12 000	1 070 176	-9 600
standard 88200 Hz	88 200	2^{21}	0.042057	24.2772	22 050	2 136 842	-17 640
standard 96000 Hz	96 000	2^{21}	0.045776	22.3453	24 000	2 140 352	-19 200
corpus (pre-#254 B2a) 96000 Hz / 2^17	96 000	2^{17}	0.73242	1.86533	24 000	174 272	-19 200
Simulation Mode interactive 48000 Hz / 2^15	48 000	2^{15}	1.4648	1.18267	12 000	54 368	-9 600

The role table for the default fifth: every grid point's musical frequency, nearest pitch and deviation, semitone class relative to the root, and the role the classifier assigns it:

k	Hz	pitch	cents	class	role
1	55	A1	0	0	Growl (sub-octave)
2	110	A2	0	0	Thickening
3	165	E3	1.96	7	Thickening
4	220	A3	0	0	Thickening
5	275	C♯4	-13.7	4	Added colour
6	330	E4	1.96	7	Thickening
7	385	G4	-31.2	10	Sourness (clash)
8	440	A4	0	0	Thickening
9	495	B4	3.91	2	Added colour
10	550	C♯5	-13.7	4	Added colour
11	605	D♯5	-48.7	6	Sourness (clash)
12	660	E5	1.96	7	Thickening
13	715	F5	40.5	8	Sourness (clash)

7.6 The result

The app's own rendering of the same measurement kind, from the user guide's figure pipeline (the bundled simulated rig fed the default fifth at the standard drive, rendered headlessly from the chart view):

7.7 Worked example and parity

The parity test runs 28 cases through both implementations, and each device case again as its phantom capture twin (9 twins): the lattice choice alone at every standard grid and both fallback grids, for the fifth, an octave, a major third and a refused pair; the polynomial devices and the two clippers at the loudest hardware drive on the shipped fifth; a phantom holding all $\sum_{k=2}^{K} 2k$ recipes at a level ladder; the identity through seeded noise at two levels and two lengths; a squarer behind a post-filter and the soft clipper behind a pre-filter; the corpus lattice with a squarer and a cubic; the wander fixture through the legacy path and through the detector; a squarer driven until its products stand above the played tones; and the worked example analyzed 100 samples late. The stimulus is the plan's own, snapped by the shipped search, so the lattice choice is

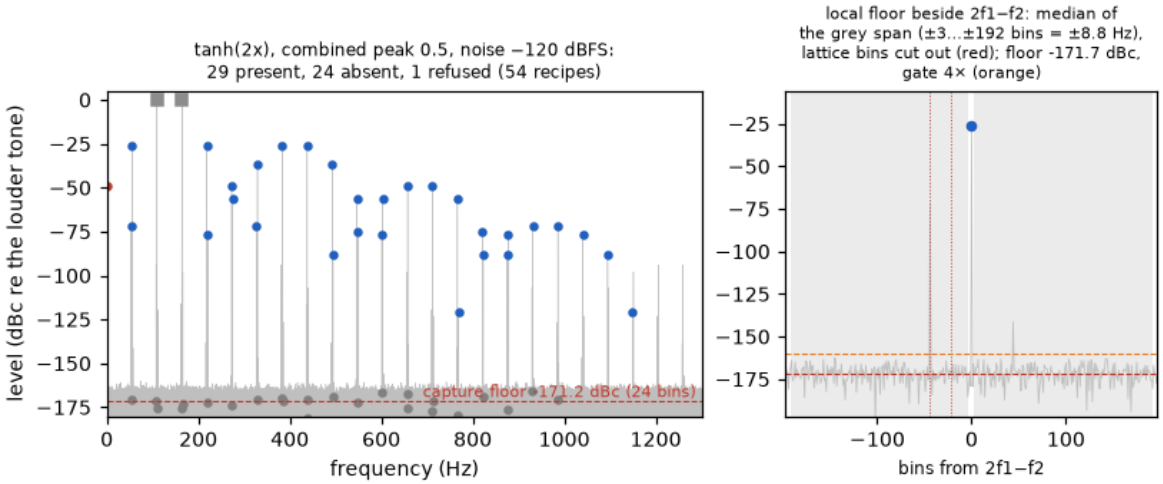


Figure 27: The worked example, $y = \tanh(2x)$ at a combined peak of 0.5 on the shipped fifth through seeded noise: the analyzed spectrum with every enumerated product marked present (blue), absent (gray) or refused (red), the two tones as squares, and the capture floor as the dashed rule (left). Right: the local-floor neighborhood of one product, the median of the gray span with every lattice bin cut out, and the gate four times above it.

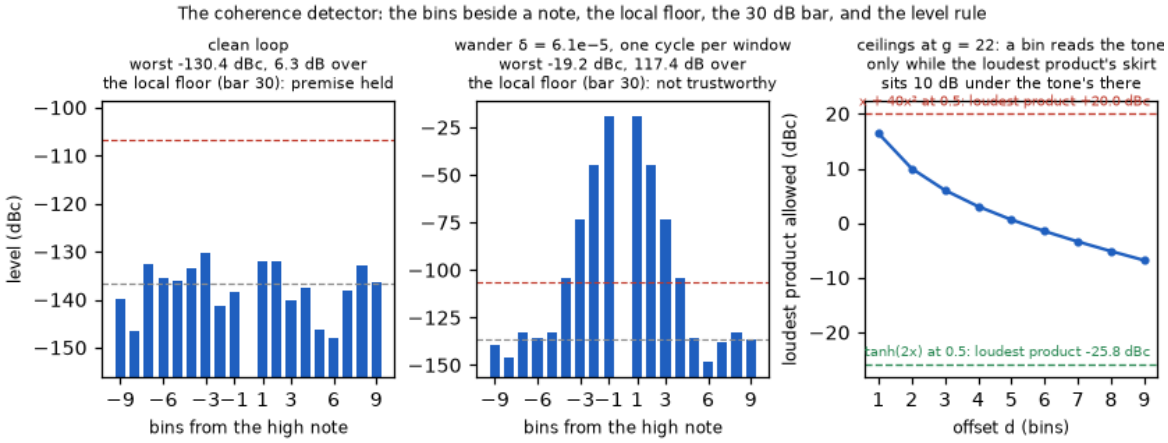


Figure 28: The coherence detector beside the high note on a clean identity (left) and on the validated wander fixture (middle), each against the local floor beside the note and the bar above it; right, the per-offset ceilings of (41) on the shipped spacing, with the worked example's loudest product far below every ceiling and a squarer driven above its own tones clearing none.

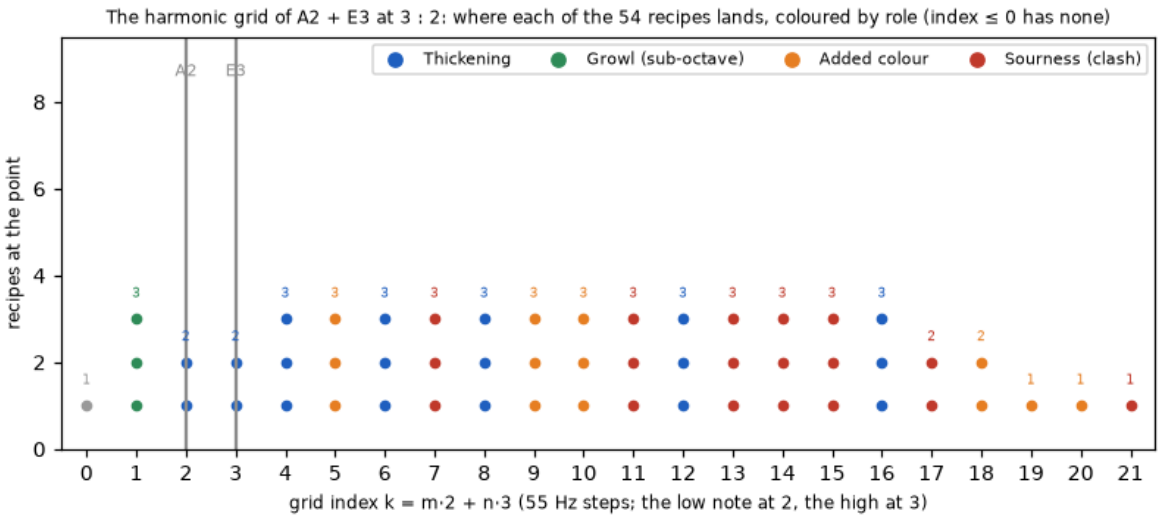


Figure 29: The harmonic grid of the requested fifth at its just ratio: every enumerated recipe at its grid index (44), stacked where several share a point and colored by the role the classifier assigns the point; the two played notes marked. Index zero and below has no role.

exercised, not assumed; the analysis is the app's own call. Beside every product the oracle prints a truth: for a memoryless device on an exact bin lattice the output is periodic on the (θ_1, θ_2) torus, so every product's amplitude is an exact two-dimensional Fourier coefficient,

$$a_{m,n} = 2\,|c_{m,n}|\,|H_{\text{post}}(mf_1 + nf_2)|, \qquad c_{m,n} = \frac{1}{4\pi^2}\iint f(A_1 \sin\theta_1 + A_2 \sin\theta_2)\, e^{-i(m\theta_1 + n\theta_2)}\, d\theta_1\, d\theta_2, \tag{45}$$

with $A_i = a_i\,|H_{\text{pre}}(f_i)|$ the drive after any pre-filter, computed by the trapezoid rule on the grid the parameter table sizes (exact for the periodic integrand's orders below the grid minus K), never by hand; the polynomial device's closed form (the binomial expansion in two sines) is printed beside the quadrature as a check on it; the tone bins carry the linear response plus the odd self-terms (a cubic's $\frac{3}{4}a_3A_1^3 + \frac{3}{2}a_3A_1A_2^2$ on tone 1); the identity's every product is zero and the phantom's are its listed levels. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. The test asserts:

- Swift against Python: every header and row quantity in decibels within 1e-07 (measured maximum 1.8e-08 on this run) on a capture with no filter in it, and within 0.001 (measured 0.00016) on one rendered through a biquad, whose direct-form recursion accumulates rounding differently in the two languages over two million samples; every discrete reading — bin, gate, reading and its evidence, grid index, role, contest, partner — identical, and every sentence identical after last-bit rounding; the two truths within 1e-08 relative (measured 2.6e-10), the closed form against the quadrature within 1e-09 (measured $\leq$ 1e-10);

- every analytically empty recipe on a noiseless capture at or below -200 dBc on both sides (measured worst -273.8): the depth at which a bin-centered rectangular window separates nothing from the loudest component is double precision's, not a window's;

- the shipped read against the truth on every content recipe, after the noise allowance (a bin holds its noise too, so a read may sit up to $20\log_{10}(1 + 10^{(\text{floor}+\text{spread}-L)/20})$ from the truth by the noise alone): within 0.001 for the bandlimited devices (measured residue at worst -1.2e-07), within 0.05 for the two kinked ones, and within 1e-08 for every phantom (measured 2.5e-10 on the full ladder);

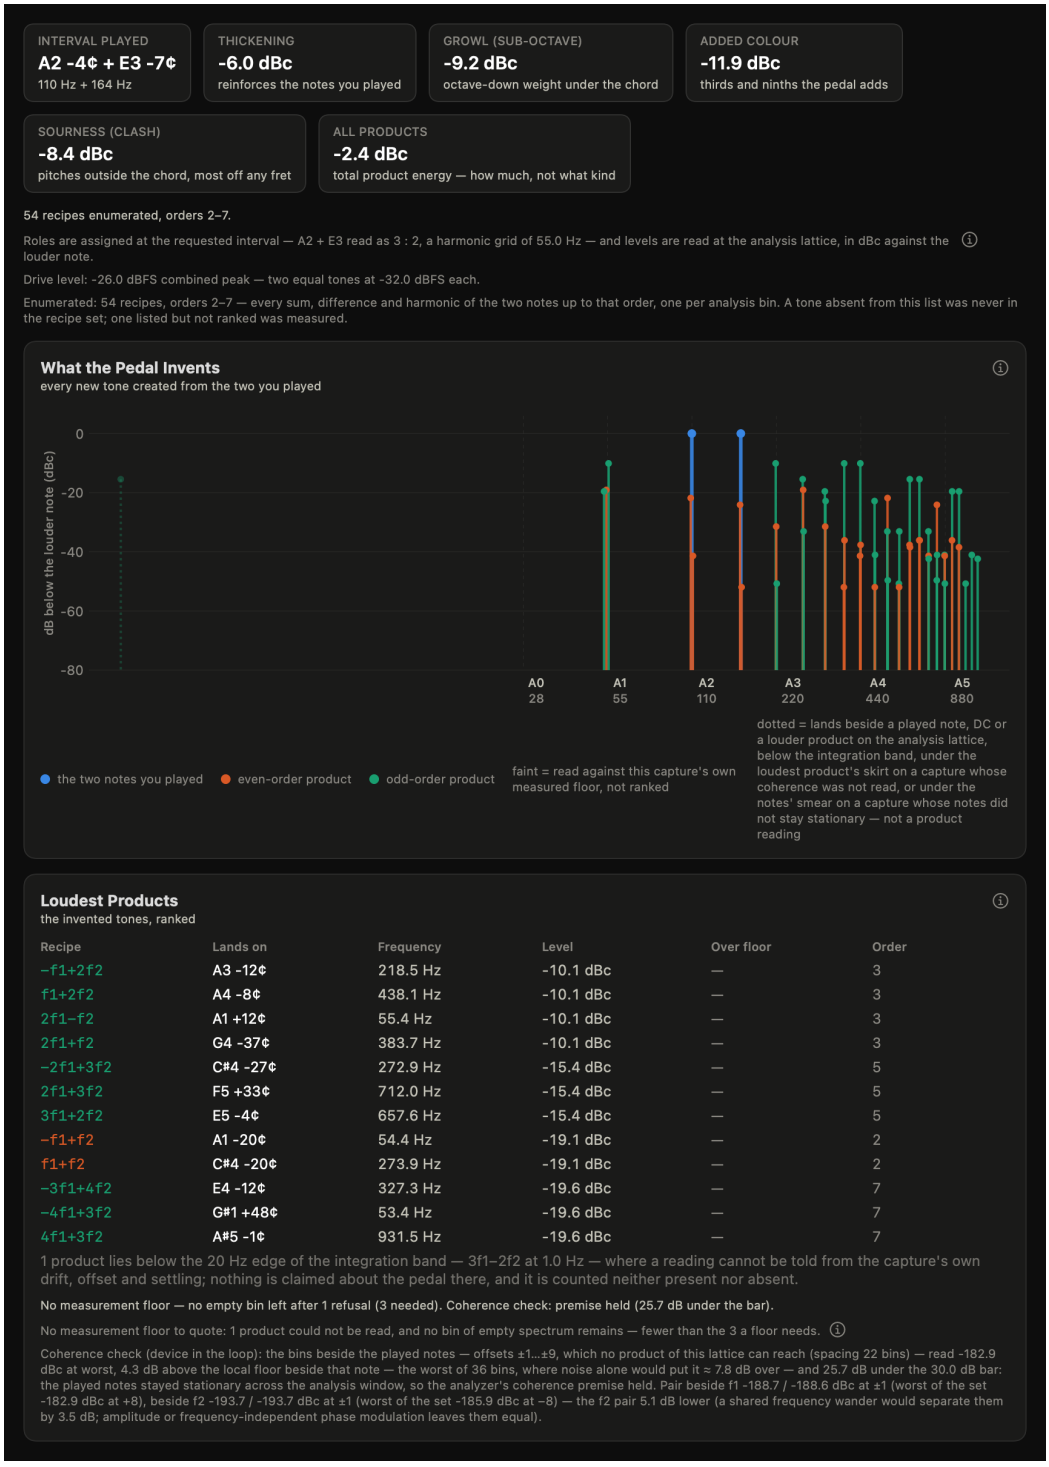


Figure 30: The Chord Roughness view as the app draws it, from the guide's simulated fixture: the pitch-axis spectrum, the role tiles in dBc, the drive line and, under the table, the floor sentence and the coherence line.

- the lattice each case played is the row the manifest recorded by running the shipped search on that grid, field for field, and the refused pair is refused with its numbers on both sides;
- the identity through noise reads every product absent, its capture floor within 3 dB of the noise's analytic per-bin median (measured deficits 0.21 and 0.14 dB at the two lengths — a median of a few dozen Rayleigh draws), and its coherence excess within 6 dB of (42) (measured 6.3 against 7.8);
- the late analysis moves the window by exactly its offset, keeps the overrun, and moves no content product's level by more than 1e-05 dB (measured 2.2e-06: the noise stream is not shifted with the window, so the late segment exchanges a hundred noise samples for a hundred new ones).

Part B's parity is established: on every predicate case — the corpus lattice's refusals by position and beside a louder product, the wander fixture through the legacy path and through the detector, and the device whose products stand above its tones — the two implementations agree on every reading, every count and both registers' sentences.

The lattice cases — the shipped fifth at the four rates, the two fallback grids, the octave at two lengths, the major third, and the refusal:

case	f_s	N	tier	$k_1 : k_2$	g	$a : b$	interval (¢)	tones (¢)
a-fifth-96k	96 000	2^{21}	standard resolvable	2398 : 3586	22	109 : 163	696.65	-3.6 / -6.94
a-fifth-48k	48 000	2^{20}	standard resolvable	2398 : 3586	22	109 : 163	696.65	-3.6 / -6.94
a-fifth-88k	88 200	2^{21}	standard resolvable	2616 : 3912	24	109 : 163	696.65	0.334 / -3.01
a-fifth-44k	44 100	2^{20}	standard resolvable	2616 : 3912	24	109 : 163	696.65	0.334 / -3.01
a-fifth-corpus	96 000	2^{17}	degraded	150 : 226	2	75 : 113	709.63	-2.15 / 7.48
a-fifth-interactive	48 000	2^{15}	degraded	75 : 113	1	75 : 113	709.63	-2.15 / 7.48
a-octave-96k	96 000	2^{21}	standard resolvable	2418 : 4823	13	186 : 371	1195.3	10.8 / 6.12
a-octave-corpus	96 000	2^{17}	degraded	150 : 302	2	75 : 151	1211.5	-2.15 / 9.35
a-third-96k	96 000	2^{21}	standard resolvable	2400 : 3024	48	50 : 63	400.11	-2.15 / -2.04
a-refusal	96 000	2^{16}	refused	25 : 50	25	1 : 2	1 200	-4.11 / -4.11

The device cases, each beside its alias-free twin. A twin is noiseless, so its empty bins are numerical dust and its "present" count is dust against dust — the verdict of a noiseless case is not compared, and the column shows why; a device case carries seeded noise far below its content. The residue is what remains of the read's error against the truth once the noise allowance is subtracted, clamped at zero:

case	content	present	S−P (dB)	floor (dBc)	vs truth (dB)	residue (dB)	at
b-poly	10	10	1.7e-09	-166.84	7.53e-06	≤ 1e-10	3f2
b-poly-twin	10	11	≤ 1e-10	-332.43	≤ 1e-10	≤ 1e-10	3f2
b-squarer	4	4	1.9e-09	-166.53	2.86e-07	≤ 1e-10	2f2
b-squarer-twin	4	6	≤ 1e-10	-331.54	≤ 1e-10	≤ 1e-10	2f1
b-asym	54	53	2.3e-09	–	0.0298	0.03	7f1
b-asym-twin	54	53	≤ 1e-10	–	1.48e-10	1.5e-10	7f1
c-tanh	30	29	1.6e-09	-171.17	0.0224	≤ 1e-10	7f1
c-tanh-twin	30	29	≤ 1e-10	-329.35	2.91e-10	2.9e-10	7f1
c-hardclip	30	29	≤ 1e-10	-156.53	0.0128	0.013	f1+4f2
c-hardclip-twin	30	30	≤ 1e-10	–	≤ 1e-10	≤ 1e-10	5f1−2f2
d-phantom	54	53	≤ 1e-10	–	2.48e-10	2.5e-10	5f1−2f2
f-squarer-hp	4	4	0.00016	-166.53	3.04e-07	≤ 1e-10	2f2
f-squarer-hp-twin	4	6	≤ 1e-10	-331.61	≤ 1e-10	≤ 1e-10	2f1
f-squarer-lp	4	4	2.1e-07	-166.53	2.09e-07	≤ 1e-10	2f2
f-squarer-lp-twin	4	6	≤ 1e-10	-331.64	≤ 1e-10	≤ 1e-10	2f1
f-tanh-prehp	30	29	4.1e-05	-171.17	0.0294	≤ 1e-10	7f1
f-tanh-prehp-twin	30	29	≤ 1e-10	-329.32	2.96e-10	3e-10	7f1
j-latency	30	29	1.6e-09	-171.16	0.0223	≤ 1e-10	7f1
j-latency-twin	30	29	≤ 1e-10	-329.35	2.91e-10	2.9e-10	7f1

Three readings from that table are worth stating in words. The two *kinked* devices — the hard clip, a near-square wave whose coefficients fall as $1/k$, and the asymmetric clipper, whose slope

breaks at zero — applied per sample at f_s fold what they make above Nyquist back onto the lattice, where a folded high-order product lands *exactly* on a low-order bin (every lattice point is a multiple of g): 0.0128 dB at f1+4f2 and 0.0298 dB at 7f1, against $\leq$1e-10 and 1.5e-10 on their phantom twins. That is chapter one's finding on a two-tone lattice, and the same corollary holds: a per-sample synthetic device is never the analytic oracle for a product near the top of the band; the phantom is. The bandlimited devices read the truth inside the allowance on every case. And the phantom of every recipe reads back 53 present and 1 out of band — the recipe the shipped fifth puts under the band's low edge, whose level still reads back to 2.5e-10 dB and whose reading is a refusal on frequency, by the band rule that outranks the gate.

The worked example, $y = \tanh(2x)$ at a combined peak of 0.5 through seeded noise: 29 present, 24 absent, a floor of -171.2 dBc from 24 bins, a coherence excess of 6.3 dB over the local floor, and a loudest product of -25.83 dBc. Its third-order difference products $2f_1 - f_2$ and $2f_2 - f_1$ read -25.832 and -25.832 dBc against truths of -25.832 and -25.832 (both present); $3f_1$ reads -37.078 against -37.078; and the second-order $f_1 + f_2$, analytically empty for an odd device, reads -173.79 dBc — the noise — and is absent, with the bin's own level as its bound.

The identity through noise: the capture floor against the analytic median $2\sigma\sqrt{\ln 2/N}$ of the same noise through the analyzer's scale, and the distance of that per-bin median below the noise's broadband rms, $10\log_{10}(N/(4\ln 2))$ — 58.79 dB at the standard length and 46.75 at the corpus length, the relation that decides whether a broadband floor can bound a bin (it cannot without an assumption about how the noise is spread, and the assumption is worth tens of decibels in the refusing direction):

case	N	noise (dBFS)	median (dBc)	floor (dBc)	bins	per-bin − broadband (dB)	excess (dB)
e-identity-80	2^{21}	-80	-106.75	-106.53	53	58.79	6.3
e-identity-110	2^{21}	-110	-136.75	-136.53	53	58.79	6.3
e-identity-corpus-80	2^{17}	-80	-94.705	-94.843	49	46.75	8.65

The predicate cases. On the corpus lattice a squarer through noise has 5 recipes refused beside the notes and DC and 8 refused beside a louder product, and its floor rests on 37 bins with those left out. The wander fixture — a shared fractional frequency wander of both tones, the validated generator of the kernel's own test, which reproduces the one real fault the reference bench has captured — reads through the legacy path at -49.25 dBc, 75.6 dB over the capture floor, with the pair difference 3.56 dB against the 3.56 a frequency wander predicts; through the detector on the shipped lattice at one cycle per window it reads -19.23 dBc, 117 dB over the local floor, over the bar in both registers. The squarer driven above its tones has a loudest product of 20 dBc, clears no ceiling, does not interpret, and marks 49 products unread under a skirt of 36.8 dB at the shipped spacing:

case	present	absent	refused	floor bins	coherence source	excess / bar (dB)	readable
g-corpus-squarer	4	37	13	37	detector	8.65 / 30	1..9
g-corpus-cubic	7	28	19	28	detector	8.65 / 30	1..9
h-wander-legacy	0	49	5	49	stored products	75.6 / 30	2
h-wander-shipped	0	49	5	49	detector	117 / 30	1..9
i-loud-squarer	4	0	50	49	detector	6.3 / 30	none

7.8 Limits and failure modes

- **The floor cannot separate the device from the rig.** A resolved product stands above this capture's noise; the measurement does not say what made it. An interface's converters intermodulate, and a bare loop resolves a handful of second-order products of its own. The capture floor is the right decision floor for presence (a per-bin read at the product's own frequency), and the wrong instrument for attribution; a control capture with nothing in the

loop is the only thing that does that job, and no rig-side broadband floor composes into a per-bin decision — the per-bin median of white noise sits 58.79 dB below the same noise's broadband rms at the standard length, so a broadband bound applied per bin would refuse genuine products by tens of decibels.

- **The percentage sums carry no floor bound.** The three capture-time sums are broadband aggregates over the products the gate passed, and a broadband quantity is exactly what an operative rig floor legitimately bounds — so a bound is available and unbuilt: a percentage summed from products that each passed their own per-bin gate can still be quoted below what the rig supports as a whole. The read-time derivation subtracts refused products and bounds nothing.

- **A coherence verdict names no culprit, and reads nothing beside a loud product.** A detector bin responds to modulation of the played tone at periods comparable to the window itself; a pedal's thermal drift, a sagging supply or a dying battery moves a note as surely as a clock, so "the notes did not stay stationary" is a statement about the capture, never about a stage. And the check is a level rule as much as a distance rule: a device whose loudest product stands above the ceiling of (41) for even the nearest offset — a full-wave rectifier's products stand above its own fundamentals — is not read at all, and the line says why. A device whose sidebands are roughly constant in absolute level and fall in dBc as its output rises is consistent with something after the level control modulating the notes; the check reports the excess, the pair shape and the level, and the design question of whether it should say more when that dependence is clean is open.

- **The smear model extrapolates, and a fast wobble escapes it.** The per-product reading of Section 7.4 extends the outer detector reads along $1/m$; a modulation faster than about two cycles per window puts its sideband beyond the anchors and the extrapolation under-reads its tail (the 2.3-cycle case still read within 1.9 dB at the margin; a faster one was not measured). The product scaling takes the larger of the two mechanisms and so over-predicts a level wobble after the clipper by up to the order ratio, which refuses more, never less.

- **What a wander looks like, and what it does not distinguish.** The pair statistic separates a shared fractional frequency wander (the pairs differ by $20\log_{10}(f_2/f_1)$, measured on the fixture to 3.56 against 3.56) from amplitude or frequency-independent phase modulation (the pairs equal); it does not separate the latter two from each other, and no verdict is drawn from the shape. A $1/d$ fall of the sidebands with offset is the rectangular skirt of a slow, non-periodic drift; a lift that does not fall with offset is broadband. A drop in a device's even-order products between a short stimulus and the current long one, alongside a skirt that follows the device position, is consistent with an operating point that shifts over the window (a rectified mean through a coupling capacitor, say); the oracle reproduces the skirt from a stated wander and does not reproduce the product drop from any mechanism it models, so that half stays a bench question.

- **Aliasing in a simulated device.** A device applied per sample folds its content above Nyquist onto the lattice, exactly on low-order bins; the kinked devices' residues above quantify it. Hardware behind a converter's anti-alias filter does not do this.

- **The enumeration is a small-signal census.** Orders two through K; at the standard drive the reference clippers' content sits in orders two and three with trace in four and five, but a hard clipper at playing levels has real energy at orders the walk does not reach. An empty row above order five means measured and nothing above the floor at this drive, never that the device makes nothing there.

- **The corpus lattice is crowded, and every record captured on it carries the marks.** At the old length the fifth's spacing was two bins: four even-order recipes sit two bins from the played notes and one two bins from DC, all refused on position, and products cluster two bins apart so the quieter of a pair is refused beside the louder. The old three-bin read applies to those records too; the separation bound was measured under it and is deliberately not raised. On an exact small ratio (the original plain snap of the fifth, an exact $2 : 3$) recipes share bins and the surviving labels are the enumeration's tie-break, marked at read time as contested; those records can never be re-measured onto a clean lattice.

- **The sub-band recipe.** On the shipped fifth one recipe, $3f_1 - 2f_2$, lands about one hertz up: under the band's low edge, where nothing a pedal does can be told from the capture's own drift, offset and settling. It reads out of band on every record — named in its own line, never in a sum, and never in a tile — and the phantom parity shows the level under it still reads back exactly.

- **The stored gate is frozen.** The presence flag and the three sums are written at capture; a change to the gate's factor or the neighborhood moves future captures only, and a read-time predicate can change which flagged bins are averaged into a floor but never which bins are flagged. The window length is likewise not stored: a consumer derives it from the stored stimulus by the plan's own rule, and a record captured at the old fixed length at a rate where it overran is asked the post-fix question and answers with no bound — never a certificate that its window was clean.

- **The interval is a few cents off, on purpose.** Keeping the products separable costs pitch inside the stated budget, disclosed on every record; a consonant interval is a simple ratio, and a simple ratio is a crowded lattice, so a request that cannot escape an exact small ratio inside the budget is refused before anything plays rather than measured on a lattice whose labels are not facts.

8 Waveform Matrix

8.1 Significance and limitations

The waveform matrix — the app's "scope view" — is the one measurement in this document with no estimator. Its payload is the capture: for every cell of a note × level grid, two sample array (encoding)s holding a fixed number of periods of the stimulus and of the device's output over the SAME time span, stored as the Float32 samples the converter delivered at the capture rate, never resampled, fitted, or cleaned up. Every other chapter's result is an estimate of a device property; a tile is the raw material those estimates are made from, and it is the picture a reader checks a model-rendered wave against. What the method decides is therefore not a number but five things about samples: WHICH notes and levels are measured (the note lattice and the amplitude lattice), what stimulus plays (chapter five's stepped tone, one per note row), WHERE each tile is cut from its row's capture (the zero-crossing start), HOW the samples are written (the encoding), and what one number is stored beside them (the row noise stamp) for a read-time judgment of the quiet tiles (the noise margin and badge).

What it cannot tell you, stated once: (a) anything about a note or a level between two lattice points — a tile is one operating point, and the lattices are deliberately sparse (a 5×5 costs 25 captures' worth of time in five rows); (b) anything a picture cannot carry — there is no harmonic level, no THD, no knee here, only the wave, and the chapters that estimate those quantities are the place to read them; (c) whether a difference between two tiles is anything but a picture — a phase shift alone redraws a wave, and the app's own sentence rides every rendering: "a visibly different wave is not automatically an audibly different one"; (d) whether a quiet tile's fuzz is the device's — the stamp says how far the tile's output stands above the row's own pre-roll, and the badge says when it does not, but the stamp is the capture's, and §8.8 says what else it can hold; (e) that a loud tile should look "more distorted" than the one below it — the Gain Map's level-trend law applies here (§6): a soft-knee overdrive's THD peaks at moderate drive and falls toward full scale as the gain compresses while the harmonic content saturates, so a loud tile that looks cleaner can be telling the truth, and the Gain Map is the cross-check; and (f) anything about what the distortion sounds like.

Every measurement described here reads the device at a constant stimulus amplitude after settling. Measurements of the time course of the device's operating point — its recovery after a burst, hysteresis between ascending and descending drive, and its response to a decaying envelope — are in development and will be documented here when they ship.

8.2 The stimulus

The two lattices. Both axes are built by refinement (nesting), never by even division, so that every smaller lattice is a subset of every larger one — $3 \subset 5 \subset 7$ — and matrices captured at different dimensions compare point for point (an even divider misses itself across dimensions: measured when the rule was chosen, its $N = 4$ lost E4 and its $N = 5$ lost G♯3 and C5). The dimensions offered are odd so that a center tile exists.

The *note lattice* over a span $[f_{\mathrm{lo}}, f_{\mathrm{hi}}]$: every candidate is SNAPPED to a note — the nearest twelve-tone equal tempered pitch with A4 = 440 Hz, except that the standard note's name returns the app's standard E2 (82.4 Hz) verbatim rather than 12-TET's 82.407 Hz (the continuity exception the guide's tuning page states) — and the snapped endpoints seed the lattice. Two guarantees are then seeded in precedence order: the note anchor (E2) joins whenever it lies strictly inside the span and $N \geq 3$, so the cross-check point exists on every lattice whose span holds the standard note; and the span's middle note, the snapped log-midpoint $\sqrt{f_{\mathrm{lo}} f_{\mathrm{hi}}}$, is the center (E4 on the default span). The remaining points are inserted per side around the center — the side with fewer interior points first, ties to the low side; within a side the widest gap in log frequency first,

ties to the lower index; a gap whose snapped midpoint collides with an endpoint is skipped — and when no gap on any side accepts a point the lattice returns fewer points honestly (a narrow span never stores one delivered note twice). On the default E2–E6 span this is the octave ladder:

N	notes
3	E2 · E4 · E6
5	E2 · E3 · E4 · E5 · E6
7	E2 · A♯2 · E3 · E4 · A♯4 · E5 · E6

The *amplitude lattice* hangs from the voltage anchor. The standard drive is a VOLTAGE at the pedal's input, $V_{\mathrm{a}} = 168$ mV peak, and dBFS is rig-relative; so on a rig whose calibration measured the volts factor V_{FS} (volts at 0 dBFS) the anchor is

$$a = 20 \log_{10}(V_{\mathrm{a}}/V_{\mathrm{FS}}) \text{ dBFS}, \tag{46}$$

and where no factor is known — the plugin render path, an uncalibrated rig — it is −26 dBFS by convention (on the reference rig, $V_{\mathrm{FS}} = 3.35$ V, the two agree by construction: -25.99 dBFS). The default floor is a floor span below the anchor, $a - 22$ dB — a constant 13.3 mV peak at the device on every calibrated rig, twelve decibels below the Reference Box P6's measured dead-zone handover — and −48 dBFS is only its no-volts expression; the ceiling stays an ABSOLUTE −12 dBFS, a rig-side constraint (the loudest standard probe rung's converter staging), which the plugin path lifts to 0 dBFS. The floor and the ceiling seed the lattice; the anchor joins when $N \geq 3$ and it lies inside the span by a stated margin; and every further point is inserted on ONE SIDE of the anchor — the side with fewer interior points, ties to the quiet side — by splitting that side's widest gap at its midpoint in dB, so the anchor sits at the center index at every odd dimension and the steps are uneven by ruling (the default 5: $-48, -37, -26, -19, -12$; 3: $-48, -26, -12$; 7: $-48, -42.5, -37, -26, -22.5, -19, -12$). With no usable anchor the lattice refines from the endpoints alone and makes no center claim; a point outside the span is never fabricated. The delivered lattices, from the manifest:

N	factor (V pk at 0 dBFS)	anchor (dBFS)	levels (dBFS)
3	none	-26	-48, -26, -12
3	3.35	-25.995	-47.995, -25.995, -12
3	4.827	-29.167	-51.167, -29.167, -12
3	4.828	-29.169	-51.169, -29.169, -12
5	none	-26	-48, -37, -26, -19, -12
5	3.35	-25.995	-47.995, -36.995, -25.995, -18.997, -12
5	4.827	-29.167	-51.167, -40.167, -29.167, -20.584, -12
5	4.828	-29.169	-51.169, -40.169, -29.169, -20.585, -12
7	none	-26	-48, -42.5, -37, -26, -22.5, -19, -12
7	3.35	-25.995	-47.995, -42.495, -36.995, -25.995, -22.496, -18.997, -12
7	4.827	-29.167	-51.167, -45.667, -40.167, -29.167, -24.876, -20.584, -12
7	4.828	-29.169	-51.169, -45.669, -40.169, -29.169, -24.877, -20.585, -12

The tie. At $N = 7$ each side's second insertion must split one of two EXACTLY equal halves of its first split (the floor-to-anchor span halved on the quiet side; the anchor-to-ceiling span halved on the loud side), and until 2026-09-27 the generator's strict comparison of floating-point widths decided it by the rounding residue of $a - 22$ and $(f + a)/2$ — 3.6×10^{-15} dB, measured. Two adjacent volts factors on one interface — 4.827 and 4.828 on a Babyface Pro, both in the library — therefore delivered different quiet-side lattices at 7 (four of seven levels apart), and one recalibration put one rig's matrices on two lattices. The tie is now broken toward the quiet side, by ruling: two gaps whose widths differ by less than 1×10^{-9} dB (the manifest's `gapTieToleranceDB`; four orders above the residue bound, seven below the 0.05 dB the value-matched comparison reads) are the same width, and the quieter is split. Both factors now deliver the quiet-gap lattice (-51.167, -45.667, -40.167, -29.167, -24.876, -20.584, -12 against -51.169, -45.669, -40.169, -29.169,

-24.877, -20.585, -12 — 0 of seven levels differ; the widest difference is the 0.002 dB their anchors differ by), and the same lattice at 5 (-51.169, -40.169, -29.169, -20.585, -12), whose one quiet-side insertion is the first split and never reaches the tie; the 3- and 5-point lattices are bit-identical to what shipped before, at every factor. Records captured before the change keep their stored lattice — a stored lattice never moves — and pair with later ones by value on the five levels every 7-point lattice of one rig shares (its 5-point lattice); the parity reproduces the resolved tie on both implementations. The fallback path with no usable anchor keeps the strict comparison deliberately: the ruling moved only future 7×7 captures, and a tolerance there would move 5-point lattices on user-moved spans; that is reported, not changed.

The stimulus. One capture per frequency ROW: a pre-roll of silence, the stepped tone at the row's note walked over the amplitude lattice quiet to loud — chapter five's stimulus exactly, with its settle region and measurement window, its cycle snap, its raised-cosine ramps, its continuous phase and its final fade, the plan passing the settle and measure durations explicitly and taking the tone's default ramp — and a tail of silence. The pre-roll is the row's own noise window, the tail lets the device's response end inside the capture. The delivered counts on the default lattice at the manifest's rate:

note	f_0 (Hz)	settle	measure	ramp	row	tile
E2	82.4	14 400	24 466	1 920	232 730	2 330
E3	164.81	14 400	23 881	1 920	229 805	1 165
E4	329.63	14 400	23 881	1 920	229 805	582
E5	659.26	14 400	24 027	1 920	230 535	291
E6	1318.5	14 400	24 027	1 920	230 535	146

Rows are captured in ascending note order; every row begins with its quietest step. The stepped tone is SHARED with Compression (and with the level-match probe): a change to it moves both instruments.

8.3 The pipeline

Given the row's capture $y[n]$, the loop's latency ℓ (the embedded calibration's round-trip estimate on the hardware path; the plugin path measures and removes its own per capture; Simulation Mode passes zero) and the stepped tone's own windows:

1. **The stamp.** Before any tile, the row's row noise stamp $\sigma = \sqrt{\sum y^2/P}$ over the capture's first P pre-roll samples — the loop plus the device's own hiss, in this capture. This is the one step of the pipeline that is not the shipped plan's: both capture sites compute it inline by the same expression, and §8.8 names it as the mirror it is.

2. **The tile's size.** $M = \mathrm{round}(p\, f_s/f_0)$ samples for p periods at the row's note — rounded to the capture rate, never resampled; the stored count is explicit in the payload.

3. **The start.** For step i with measurement window $[w_{\mathrm{lo}}, w_{\mathrm{hi}})$ in the tone's own timeline, the zero-crossing start is the first sample at or after the window's start whose stimulus phase is a whole number of cycles:

$$s_i = \max\Big(\mathrm{round}\big(\lceil w_{\mathrm{lo}} f_0/f_s \rceil\, f_s/f_0\big),\ w_{\mathrm{lo}}\Big), \tag{47}$$

   so every tile begins where the input rises through zero. A start-index choice, not a resample.

4. **The window check.** If $s_i + M$ runs past $w_{\mathrm{hi}} - R$ (the window's end less the ramp) the plan THROWS — a tile is never cut across the next step's ramp. With the default period

count this cannot fire at any note (the measurement window holds at least four cycles and a tile two); it fires when the period count exceeds the window's cycles, and §8.7 runs it.

5. **The two tiles.** The input tile is the tone's samples $s_i \dots s_i + M - 1$ as Float32; the output tile is the capture's samples $P + \ell + s_i \dots P + \ell + s_i + M - 1$ as Float32 — the SAME time span at the latency, so the device's phase shift between input and output is preserved in the picture, never aligned away. If that span runs past the capture the plan THROWS; a silently short tile would be a fabricated measurement.

6. **The payload.** Each tile's two arrays carry their encoding name (`float32le-base64`), their channel and their sample count explicitly, and the result carries the ACTUAL delivered frequencies and levels — never the requested ranges — the rate, the period count, the tiles in row-major order and, as an additive field, the row stamps. A reader that finds another encoding, or a byte count that disagrees with the stated sample count, refuses rather than guesses.

8.4 Read-time facts on the result

Everything below is computed from the frozen payload when a record is read; nothing derived is stored.

Statistics and the badge. A tile's decoded output has an rms, a peak (the tile is DRAWN scaled to its own peak, and the pane says so) and a mean (the evidence §8.8 reads). The noise margin and badge's margin is $20 \log_{10}(\text{rms}/\sigma)$ against the row's stamp — nil when the record carries no stamp (absence is honest: no badge either way), $-\infty$ for a silent tile — and a tile whose margin is under the clear margin (the parameter table gives it) wears the "$\approx$ noise" badge, whose sentence says most of the drawn wave is the rig's and device's own noise and never that the tile is below a floor, because the wave is drawn.

The cross-check. Because the note anchor is on every lattice that spans it and the amplitude anchor is on every lattice, the tile at (E2, the standard drive) is the cross-check point: the operating point at which the Transfer Curve's measured cycle (§4) was captured, by construction, whenever that record uses its default note. §8.7 measures the agreement once.

APP LEVEL. Four rules live in the app and not in the kernel; the oracle cannot call them, so they are stated and census-listed, never paritied. The pane states the anchor sentence only when the record's own lattice carries the rig's anchor within a stated tolerance, and the ceiling clause only when the loudest stored level is the default ceiling (evidence-gated on the stored values; a user-moved span gets nothing). Every tile is drawn scaled to its own peak, and the sentence beside the grid says so, with the app's own disclaimer that a visibly different wave is not automatically an audibly different one. The Compare view's value alignment (app level) matches two records' tiles on the STORED delivered notes and levels — notes within a stated ratio (about 1.7 cents), levels within a stated tolerance — never on row or column position, the level basis decided by the same predicate that draws the axes: matched on dBFS where the two records share an axis, on the DELIVERED VOLTAGE across two calibrated rigs with different factors (the floor and the anchor are voltages and pair; the rig-side ceiling generally does not), and not at all across a plugin/hardware pair, which is stated and falls back to two solo grids; three or more records take the intersection across all of them, the first record's values the match reference. And the alignment sentence is always shown: a silent alignment is the failure.

8.5 Parameters

Every value below is read from the manifest, which read it from the shipped code: every static by name, the lattices by calling the shipped generators at every dimension and at four volts factors, the durations off a plan instance, the delivered per-row counts off the stepped tone, the note names by the shipped pitch code. The anchor's admission margin and match tolerance, the balance rule's ties, the rounding in the tile count and the start are literals the manifest carries as rule strings.

Parameter	Value
Lattices	
dimensions offered / the default	3, 5, 7 / 5
default note span (low / high)	E2 (82.4 Hz) / E6 (1318.5 Hz)
note anchor (the transfer curve's standard note)	E2 (82.4 Hz) — equal to the transfer default
voltage anchor / its no-volts expression	168 mV pk / -26 dBFS
default floor: span below the anchor / no-volts expression	22 dB / -48 dBFS
ceiling (absolute) / plugin ceiling	-12 / 0 dBFS
anchor admission margin / anchor match (literals)	0.5 dB / 1e-12
Stimulus	
pre-roll / tail silence	0.2 s / 0.2 s
settle / measure (passed by the plan) / ramp (the tone's default)	0.15 s / 0.25 s / 0.02 s
periods per tile (the default)	2
pre-roll / tail samples at 96 000 Hz	19 200 / 19 200
Encoding	
sample array encoding	`float32le-base64`
Noise	
clear margin (the badge fires below it)	6 dB over the row stamp
App level (stated, not paritied)	
anchor disclosure tolerance	0.05 dB
compare: level tolerance / note tolerance	0.05 dB / 0.001 (≈ 1.7 cents)
Oracle (`analysisdump matrix-synth`)	
crossover dead zone (the kernel test's)	0.016406 (-35.7 dBFS)

8.6 The result

The guide's page for this kind carries no rendered figure, so the chapter includes no app rendering; every figure above is the reimplementation's own.

8.7 Worked example and parity

The parity test runs 26 cases through both implementations. The Python imports chapter one's device family and noise generator, the transfer chapter's filters and steady-state series, and chapter five's stepped tone, and owns the rest: the two lattice generators, the anchor, the row stimulus, the cut, the Float32 rounding, the statistics, the stamp, the margin and the badge, the crossover device, and the two ways of rendering the rows. Every case is a known device through a known LOOP — a delay, one seeded noise stream per row over the whole capture, a chain gain — over the shipped lattices, and beside every tile the oracle prints its TRUTH: for a memoryless device with no filter, the device's function of the float64 tone at the tile's own samples, exact; for a

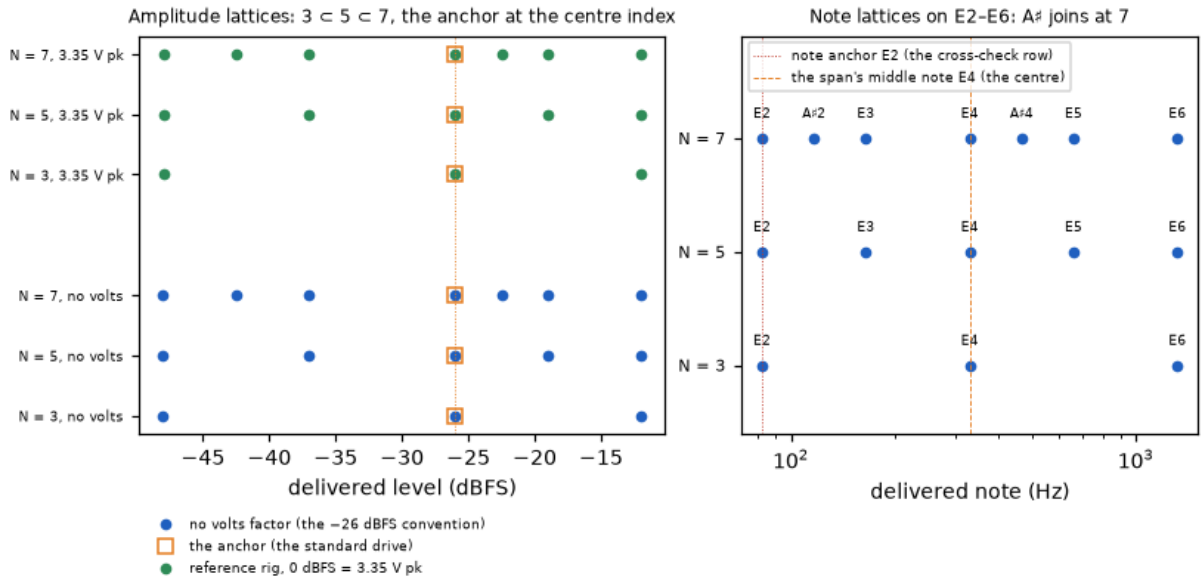


Figure 31: The two lattices at $N = 3, 5, 7$. Left: the amplitude lattice on the no-volts convention and on the reference rig's factor (equal by construction), the anchor at the center index of every dimension, the floor a fixed span below it, the ceiling fixed; every smaller lattice is a subset of the larger. Right: the note lattice on the default span, the octave ladder at 5 with A♯2 and A♯4 joining at 7; E2 (dotted) is the cross-check row and E4 (dashed) the center — two different rows.

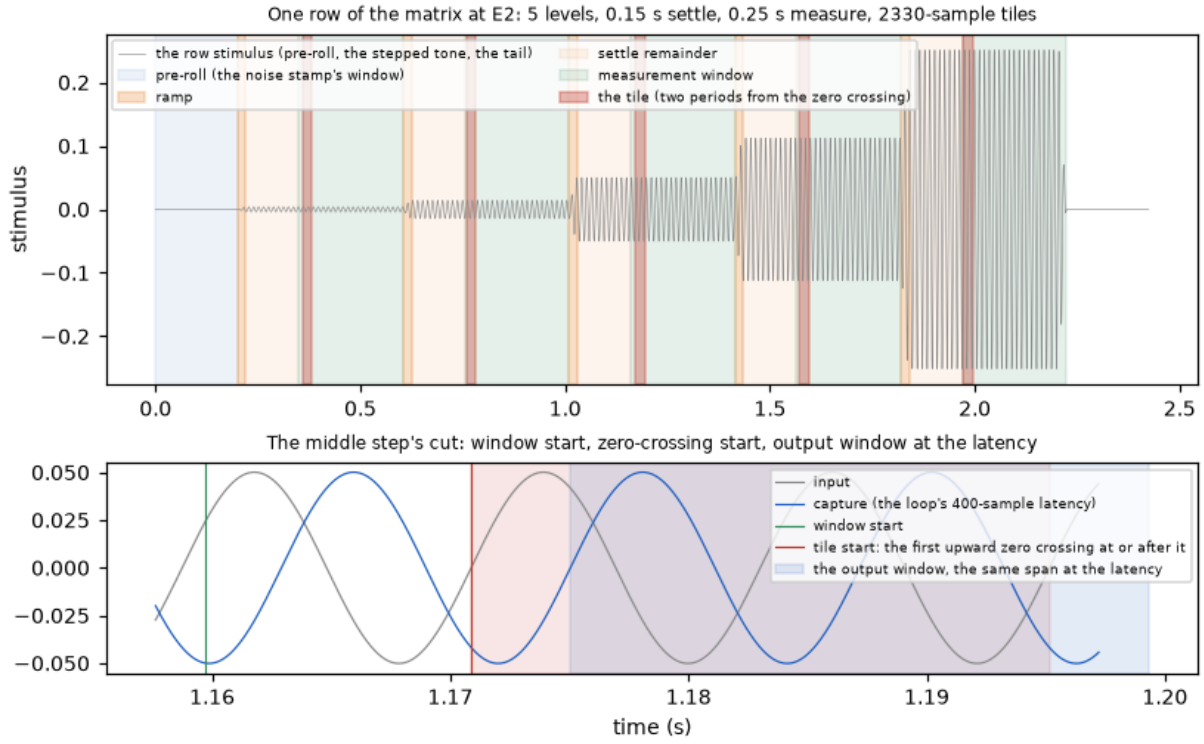


Figure 32: One row's timeline at E2 on the default lattice (top): the pre-roll, each step's ramp, settle remainder and measurement window, and the two-period tile cut from each window. Below, the middle step's cut: the window start, the first upward zero crossing at or after it where the tile starts, and the output window — the same span at a 400-sample loop latency, so the capture's own delay shows as a shift rather than being aligned away.

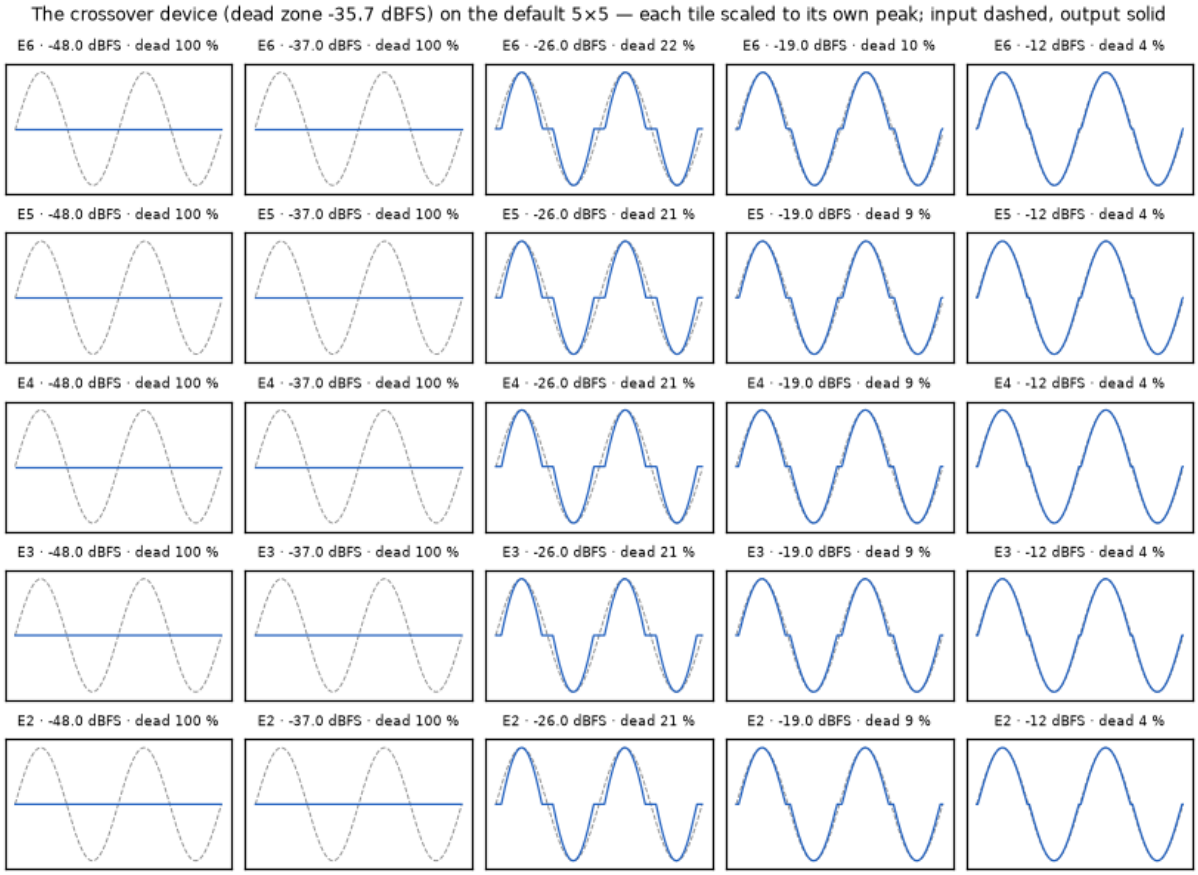


Figure 33: The crossover device (the kernel requirement test's dead zone, the Reference Box P6's measured handover) on the default 5×5, drawn as the app draws a tile: input dashed, output solid, each tile scaled to its own peak — so shapes are read in the tiles and levels on the labels, and a visibly different wave is not automatically an audibly different one. The two quietest columns lie entirely inside the dead zone (a flat line at 100 %); at the anchor a fifth of each cycle is dead; at the ceiling a sliver — the physics the default floor was set to reach.

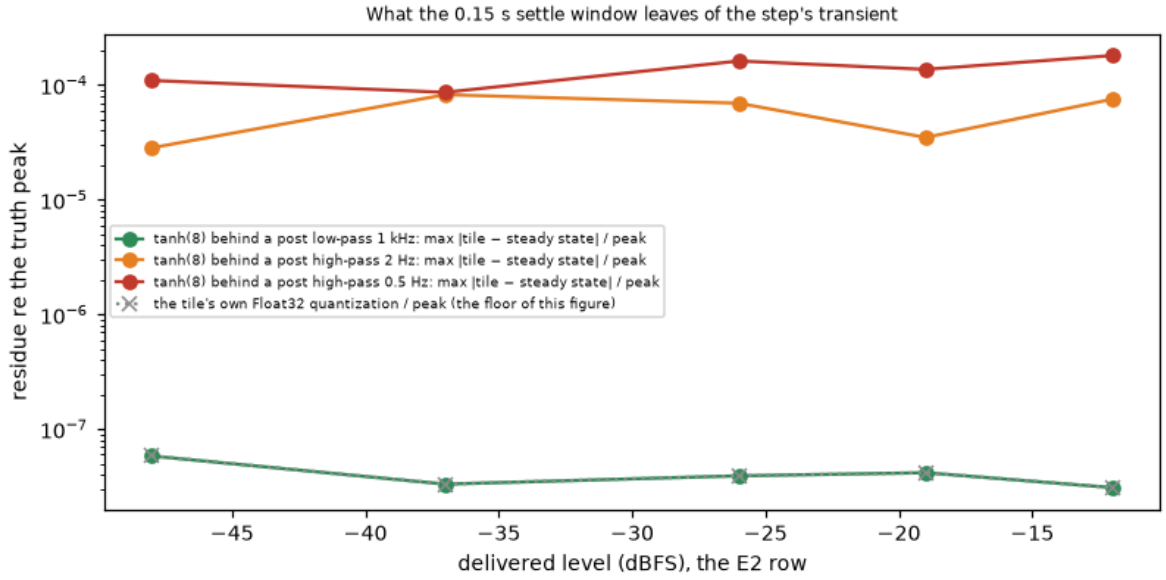


Figure 34: What the settle window leaves of a step's transient: the stored tile's maximum departure from the steady-state series, relative to the truth peak, per level of the E2 row, for $\tanh(8x)$ behind a 1 kHz low-pass (absorbed — the residue is the tile's own Float32 quantization, drawn as the floor) and behind a 2 Hz and a 0.5 Hz second-order high-pass, an output coupling capacitor whose time constant is of the settle's order and longer.

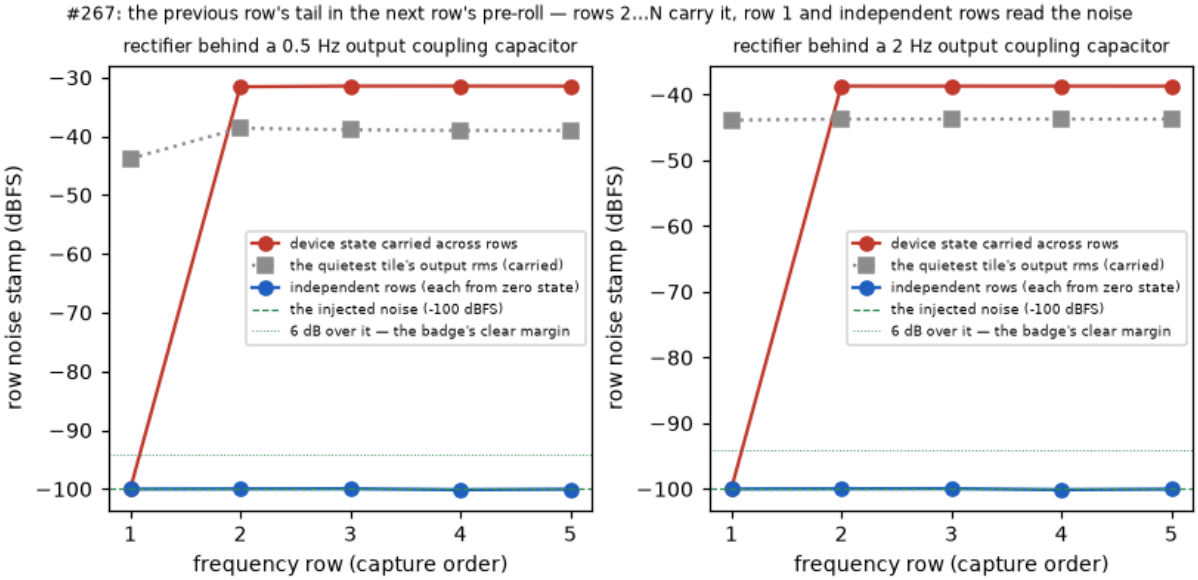


Figure 35: The stamp's limit, reproduced: a rectifying device (an asymmetric clipper) behind an output coupling capacitor, with white noise injected at −100 dBFS. Under carried-state rendering (the hardware truth: one device state across the rows) rows 2 to 5's pre-rolls hold the previous row's loudest step decaying through the capacitor and the stamp reads tens of decibels over the noise, while row 1 and every independently rendered row (Simulation Mode's shape) read the noise itself; the quietest tile's own rms sits UNDER the contaminated stamp.

device behind filters, the steady-state series through the filters' analytic responses (the transfer chapter's, at each sample's stimulus phase); for the stamp, the injected noise. The oracle renders the rows through ONE device state in capture order by default (carried state — a coupling capacitor does not reset between the rows of one matrix) and, on request, each row from zero state as Simulation Mode renders them.

Where the method is exact the parity's bar is EQUALITY, and the provenance says why. A measured figure under 1e-10 is printed as the bound $\leq$ 1e-10: below it the digits are the machine's last-bit arithmetic rather than the method's, and two machines running the same code do not share them. A stated tolerance keeps its digits. The test asserts:

- Swift against Python, EXACTLY: every delivered lattice at every dimension and factor (the tie included), every note name and cents figure, every start index in the tone and in the capture, every sample count, every badge, the truth's kind, every dead fraction — 0 mismatches on this run; the encoded payload's size to 8 bytes (the measured difference is a stamp's shortest spelling differing in its last digit — a machine's last bit, which the fragment does not print);
- Swift against Python on every statistic: the same Float32 samples summed in different orders — within 1e-12 relative on an unfiltered device (measured $\leq$ 1e-10) and 1e-07 through a biquad (measured 1.2e-08, a slow recursion rounding a handful of samples across a Float32 ulp); the means within 1e-08 (a near-zero mean has no relative scale; measured 6.1e-10); the error quantities within 1e-07 (measured 3.7e-09); the margins within 1e-05 dB (measured 1.7e-07); the stamp within 1e-07 relative with noise injected (measured 8.2e-09); the cross-check's figures within 1e-12 (measured $\leq$ 1e-10);
- the stored tile against its truth on 14 memoryless cases, EQUALITY to Float32 rounding: on every tile the departure from the exact truth equals the tile's own quantization, $\max |\text{Float32}(y) - y|$ (measured $\leq$ 1e-10 beyond it; the quantization itself 5.9e-08 relative to the peak, under a Float32 ulp of 5.96e-08);
- the settle window: behind a fast filter the residue after the quantization is $\leq$ 1e-10 (the transient absorbed; bar 1e-10), the relative residue 5.8e-08 (bar 1e-07); behind a 2 Hz and a 0.5 Hz coupling capacitor the residue is what the window LEFT — 8.3e-05 and 0.0013 of the truth peak (-81.7 and -57.6 dB), pinned above the quantization and under 0.005;
- the stamp against the injected noise on the identity: -0.12 to 0.0617 dB over five rows (the pre-roll estimate's own scatter over 19 200 samples; bar 0.3); the badges: none on the identity; the same device and noise through a loop 12 dB louder: every lattice value, start index and badge equal, every output statistic and the stamp moved by 12 dB with a departure of 4e-08 (bar 1e-06); a capture 100 samples late with the tiles cut at 100: every stored statistic identical ($\leq$ 1e-10); a window cut 1 sample late: the residue against the truth 1.12 % of the peak at E2 and 17.9 % at E6 — pinned in order only, since a one-sample phase error grows with the note and on a flattened wave it moves the edges:

note	E2	E3	E4	E5	E6
residue (% of peak)	1.12	2.25	4.49	8.98	17.9

- the crossover device: the quietest tile entirely inside the dead zone (a dead fraction of 1), the loudest mostly outside (0.0412), the fraction monotone down the level axis (1, 1, 0.212, 0.0936, 0.0412 on the E2 row) and equal to the truth's, on every row and on both sides — the kernel's founding requirement through the whole pipeline;

- the four refusals, run on the real code and reproduced: an unknown encoding and a count mismatch decode to nothing; a capture one sample short of the last tile throws, with its message; a period count no window can hold throws, with its message;
- the payload: a 7×7 on the default span at 96 kHz encodes to 506 597 bytes (a $5 \times 5\,248\,648$), inside the kernel's own bound.

Part B's parity is established: the delivered lattices at every dimension and factor (the tie included), every start index, the two decoder refusals and the two extraction refusals with their messages, the stamp, every margin and every badge agree between the two implementations on every case.

Every case, with its exact-mismatch count, its worst statistic difference, its worst error-quantity difference, its quantization, its residue beyond the quantization and its worst residue relative to the truth peak:

case	kind	exact	S−P stat	S−P error	quantization	truth residue	re peak
a-identity	memoryless	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
b-hardclip	memoryless	0	≤ 1e-10	≤ 1e-10	3.7e-09	≤ 1e-10	5.8e-08
b-tanh8	memoryless	0	≤ 1e-10	≤ 1e-10	3e-08	≤ 1e-10	5.8e-08
c-crossover	crossover	0	≤ 1e-10	≤ 1e-10	7.4e-09	≤ 1e-10	5.5e-08
d-tanh8-prehp	fast	0	≤ 1e-10	≤ 1e-10	3e-08	≤ 1e-10	5.8e-08
d-tanh8-postlp	fast	0	≤ 1e-10	≤ 1e-10	3e-08	≤ 1e-10	5.8e-08
d-tanh8-posthp2	slow	0	2.5e-10	1.6e-10	3e-08	7.3e-05	8.3e-05
d-tanh8-posthp05	slow	0	3.3e-09	9.4e-10	3e-08	0.00017	0.0013
e-3-none	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.8e-08
e-7-none	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
e-5-335	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
e-7-335	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
e-7-4827	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
e-7-4828	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
e-5-4828	lattice	0	≤ 1e-10	≤ 1e-10	1.5e-08	≤ 1e-10	5.9e-08
f-hardclip-noise	gain	0	≤ 1e-10	≤ 1e-10	3.7e-09	4.1e-05	0.0099
f-hardclip-gain12	gain	0	≤ 1e-10	≤ 1e-10	1.5e-08	0.00017	0.0099
g-latency100	latency	0	≤ 1e-10	≤ 1e-10	3e-08	≤ 1e-10	5.8e-08
g-miscut1	miscut	0	≤ 1e-10	≤ 1e-10	3e-08	0.17	0.18
h-identity-noise	noise	0	≤ 1e-10	≤ 1e-10	1.5e-08	4.1e-05	0.0099
i-asym-hp05	rectifier	0	1.2e-08	2.3e-10	3e-08	0.018	1
i-asym-hp05-indep	rectifier	0	5.4e-10	2.6e-10	3e-08	0.018	0.092
i-asym-hp2	rectifier	0	3.3e-10	≤ 1e-10	3e-08	0.011	0.13
i-asym-hp2-indep	rectifier	0	1.3e-10	3.7e-09	3e-08	0.011	0.042
k-plugin	plugin	0	≤ 1e-10	≤ 1e-10	3.7e-09	≤ 1e-10	5.8e-08
k-floor40	floor	0	≤ 1e-10	≤ 1e-10	3.7e-09	≤ 1e-10	4.7e-08

Two findings the oracle settles. The first is the stamp's. Under carried state a rectifier behind a 0.5 Hz output coupling capacitor reads its first row's stamp 0.0074 dB from the injected noise and rows 2 to 5 68.4 to 68.6 dB OVER it — a factor of 2 641 — while the same device on independent rows reads every row within 0.12 dB of the noise; behind a 2 Hz capacitor the carried rows read 61.3 dB over and the independent ones within 0.12. The library's own instance (the Reference Box's two rectifying positions on two dates, rows 2 to 7 at 8 to 16 times the first row's stamp on a 30 kΩ load and 100 to 230 times on a 1 MΩ one) is the same mechanism: the previous row's loudest step decaying through the capacitor into the next row's pre-roll — a DC tail, not noise. The ratio is set by the device's DC shift, the capacitor's time constant and the rig's noise, so the oracle's factor is not the library's and is not meant to be; the shape is — rows 2 to N only, and nothing at all when the rows are rendered independently, which is why Simulation Mode cannot show it. The quiet tiles' residue against the steady state tells the same story from the tile's side: 1.02 of the truth peak on the carried rows against 0.0919 on independent rows, the tail reaching the tiles. The second is the cross-check. The identity's tile at (E2, the standard drive), read at the transfer analyzer's bin-center phases, differs from the transfer curve's measured cycle by 1.18 % of the cycle's peak (-38.5 dB; the test pins the control under 0.02 of the

peak) — and the identity has nothing to differ by, so that is the FRAME: the transfer analyzer averages every cycle after its skipped start cycles, the tone's end fade included, and its cycle sits under the tone by the fade's share of the average (an input-cycle deficit of -0.0324 dB; the anchor tile's amplitude, $10^{-26/20}$, is also 0.0206 dB over the transfer plan's 0.05). The worked example, $\tanh(8x)$, differs by 1.15 % (-38.8 dB) — the same size, so at the cross-check point the two instruments agree to their frames' own difference, measured once and stated as the "same operating point by construction" claim's actual content.

8.8 Limits and failure modes

- **The stamp reads whatever the pre-roll holds.** A pre-roll rms is a noise figure only when the pre-roll holds noise. A rectifying device with an output coupling capacitor puts the previous row's loudest step into the next row's pre-roll as a decaying offset, and the stamp reads the tail — reproduced above, measured in the library on two dates. The badge then reads against a floor too high, and a device with a longer tail or a lower output would wear "≈ noise" on its quiet tiles for a DC offset. Whether the stamp should be mean-removed, the pre-roll lengthened, or the DC stated beside the rms is undecided and filed; nothing capture-time moved here.
- **The stamp is a mirror.** The row noise stamp is not in the plan: the app's measurement service and Simulation Mode's measurer each compute it inline by one expression, and the oracle reproduces that one expression and says so. A plan-owned stamp — the same follow-up that gives Compression its plan — retires the mirror; until it lands, the stamp is the one step of this chapter's pipeline that is not the shipped plan called.
- **The settle window is a duration, and a coupling capacitor can be slower.** The tile is cut from the settled part of a step, and behind a fast filter the settle absorbs the transient to rounding; a slow output capacitor leaves a residue the parity measures, and a rectifier's DC step at every level change decays through the whole tile. The residue is stated, never removed: the tile is the capture.
- **A one-sample latency error moves the output tile by one sample.** The output window is placed at the calibrated latency, and a wrong latency shifts it; the residue grows with the note (the same sample is a larger phase error at E6 than at E2) and on a flattened wave it moves the edges. The latency is the calibration's, measured once per rig state, and the picture preserves the device's phase shift precisely because it is not re-aligned.
- **The tie.** Two adjacent volts factors delivered different quiet-side 7-point lattices until 2026-09-27; both were valid, and the value-matched Compare pairs what they share. The tie is broken toward the quiet side now, so a 7-point matrix captured before the change shares five of its seven levels with one captured after on the same rig — the picture is the same at those five, and the pane says how many tiles line up.
- **The transfer cross-check reads the frame.** At the cross-check point the two instruments agree to the transfer cycle's own end-fade deficit; a reader comparing the tile to the transfer pane by eye is comparing against a cycle a fraction of a percent low.
- **One operating point per tile, and the lattices are sparse.** Nothing between two points is measured, and a 3×3 has no point at all in the 22 dB gap between the floor and the standard drive — which is why the default is 5×5.
- **No app rendering.** The guide page has none, so the chapter shows none; the figures are the reimplementation's.

9 Glossary

Generated from `Docs/Methods/terms.yaml`; the plain-language companion prints the same entries in plain language.

alias image Content above half the sample rate folds back below it in a sampled capture; order $k+1$'s image collides with the read of order k once $f_0 > f_s/(2k+1)$.

amplitude lattice The matrix's level columns in dBFS: the floor and the ceiling seed, the rig's anchor joins when it lies inside the span, then anchor-centered per-side refinement splits the widest gap on the emptier side so the anchor sits at the center index at every odd dimension and the lattices nest.

analysis lattice The set of bins the two tones and every product occupy: with tone bins $k_1 = ga$ and $k_2 = gb$ ($a : b$ coprime), the product (m, n) lands on bin $mk_1 + nk_2$ exactly, so every lattice point is a multiple of g and the closest two products can come is g bins. Chosen once by the ratio search and scaled to each grid.

analysis window (N) Exactly N captured samples, rectangularly windowed, starting after the fade-in and the settle and ending before the fade-out; N is the longest power of two under the ceiling whose segment fits the steady state, and the stimulus is sized to hold it — 2^{21} samples (21.8 s) at 96 kHz.

analyzer residue The THD the shipped analyzer reads with nothing in the loop — the plan's own sweep handed back as its capture — per column of the record's grid, the worst over the sweep lengths the sheet offers, computed at read time from the record's geometry. One column above the sweep's start the fade-in ramp leaks coherently into the second-harmonic window, and a noise floor cannot see a leak that scales with the signal.

audition model The per-level Hammerstein fit the record carries beside the grid, consumed by the Hear-it panels only: no measured quantity reads it, and a render between two of its levels invents a bracket-shaped clipping the device does not have.

averaged cycle The stimulus and the capture folded onto one period of the drive note in n equal phase bins, each bin holding the mean of every sample whose phase fell in it, from the skipped opening cycles to the end of the tone; the figure's two coordinates are the two bins' means.

bound A reported quantity stated as an inequality rather than a value: a read at or below its applicable floor is quoted as at most the floor (an upper bound on what the device could have contributed), a role sum with an unread member as at least its sum, a knee inside the floor-bound zone as at or below that level. The honest claim where the data cannot fix a point; never a figure.

box normalization Every loop area is divided by $4\,x_{\max}\,y_{\max}$, the rectangle that bounds the figure, so the areas are dimensionless and a pure ellipse from a phase shift φ reads $(\pi/4)\,|\sin\varphi|$.

calibration stamp The embedded calibration's THD-versus-frequency curve: one read of the loopback's own distortion-plus-noise at each of a log-spaced set of fundamentals, measured once at the calibration's own drive, stored in every record.

capture floor The read-time summary of a capture's empty bins: the upper median of the levels of every product judged not present and not refused on lattice position or band, with the spread above it. It bounds what the capture could see and can change no product's verdict — one direction, no feedback.

carried state In the oracle, rendering every row of a matrix through one continuous device state in capture order — a physical pedal's coupling capacitor does not reset between rows — so a previous row's tail can reach the next row's pre-roll; the alternative, each row from zero state, is Simulation Mode's shape.

cell One (level, column) of the summary grid: a THD figure over the declared band and one figure per stored harmonic order, each read from that level's own analysis at the column's fundamental.

center tile The middle cell of every matrix: the standard drive (the amplitude anchor at the center index) at the span's middle note (E4 on the default span) — the center guarantee, distinct from the cross-check guarantee and never conflated with it.

cleanup reading A THD-versus-level curve read as one of three states: the interpolated level at which it first reaches the threshold (cleans up); at or above the threshold already at the quietest level (never clean, within the probed range); under it at every level (always clean). On a Gain Map it is a column's read and the tile reads the column nearest the Compression probe's note; on a Compression record it is the curve's own read, the same three states by the same rule, the top of the probed range never dressed up as a crossing.

clipping A memoryless nonlinearity that limits the output amplitude: the output follows the input up to a threshold and is compressed or held beyond it (hard: a corner into a flat ceiling; soft: a gradual bend). An odd-symmetric clipper produces odd harmonics only; any asymmetry between the two half-cycles adds even harmonics.

coherence (Γ) $|\sum_i z_i| / \sum_i |z_i|$ over lagged products $z_i = H(b_i + \text{lag})\,\overline{H(b_i)}$: shaped content is a phasor whose phase advances smoothly with frequency, noise scatters isotropically, and the statistic is $1/\sqrt{n}$ for pure noise.

coherence check Whether the played tones stayed stationary across the window: the worst of the detector bins beside each tone, read against the local floor beside that tone, against a measured 30 dB bar — interpreted only over the offsets where the tone's own skirt outweighs the loudest product's, and only as a verdict about the chain, never a culprit.

coherence detector The raw spectrum at offsets $\pm 1 \ldots \pm 9$ bins beside each played tone, stored at capture with the local floor beside the tone; the widest offset is derived from the narrowest resolvable spacing ($13 - 1 - 3$), and every offset is read because an even-cycle wander lights the even ones only.

coherent sampling Both tones complete a whole number of cycles in the analysis window, so a rectangular window leaks nothing: every product sits on one bin and the floor beneath a record is noise, not a window's skirt. The premise has no margin — a tone that drifts during the capture spills a $1/(\pi m)$ skirt across the lattice.

composed margin A cell's THD in dB over the applicable floor, the floor being the larger of the rig-noise bound (the median-read stamp with its drive term, or the operative floor) and the analyzer residue — the larger floor is the smaller margin. Nil where neither rig source is usable.

compression curve The touch-response measurement (§6.3): one tone stepped through a lattice of input levels, and at each step the fundamental's output amplitude and the THD at the probe note — output level against input level, both in dB, with THD against level as its companion.

corroborated peak A cell that stands clear of its floor by the minimum margin AND whose neighbor one column either side — the same level, a third of an octave away — also stands

clear; a real peak has a shoulder in frequency, a floor-riding read has none. Corroboration is never asked for in level.

coupling capacitor A series capacitor at a device's output: a first-order high-pass whose corner sits far below the probe note, so it removes the output's mean and leaves every harmonic's magnitude essentially untouched — an asymmetric clipper's unequal clip levels are centered by it while its even harmonics survive.

cross-check point The tile at the note anchor and the amplitude anchor — the operating point at which the transfer curve's measured cycle was captured by construction, whenever the transfer record uses its default note; on the default span a corner row, never the center.

cycle snap The measurement window's length rounded to a whole number of periods of the probe note (at least four), so the window holds an integer number of cycles of the fundamental and of every harmonic, and the one-bin estimator sees no leakage from the tone itself.

dBc Decibels relative to a carrier — here the fundamental, or the louder of two played tones — so that a product's level is stated as a ratio to the signal that produced it.

dBFS Decibels relative to digital full scale: $20\log_{10}$ of a linear amplitude in which 1.0 is the largest value the converter or file can represent.

DC tail The decaying offset an output coupling capacitor holds after a step with a non-zero mean ends: on a rectifying device the previous row's loudest step is still discharging through the capacitor into the next row's pre-roll, where the row noise stamp reads it as if it were noise; a deterministic one-way decay, not noise, set by the device's DC shift, the capacitor's time constant and the load.

decibel (dB) A logarithmic ratio: $20\log_{10}$ of an amplitude ratio (equivalently $10\log_{10}$ of a power ratio). $+6$ dB doubles an amplitude, -20 dB is one tenth, -40 dB one hundredth; the reference is named by a suffix or by context.

deconvolution Dividing the captured response's spectrum by the sweep's spectrum — here, multiplying by the analytic inverse filter — to recover the device's impulse response, in which the harmonic pulses appear at their time advances.

delivered fundamental A cell's drive level plus its stored H_1, in absolute dBFS. The journey's harmonic cells are uncompensated — the loop's gain is in them — so this is the level the fundamental actually reached at the converter, the quantity the floor's drive term compares against.

derotation Multiplying every harmonic response by $e^{+j2\pi f\delta/f_s}$ on its own output-frequency axis to remove a capture delay δ — a unity-magnitude rotation, so magnitudes are untouched at every bin.

drive level The digital level of one level's sweep, $20\log_{10} A$ in dBFS, where A is the sweep amplitude the device is driven with. On a hardware record it is the interface's output level; on a plugin record it is digital full scale itself. The lattice is a list of these.

drive term The floor's adjustment for a cell delivered quieter than the calibration: the stamp's ratio rises by the difference, because the rig's noise is absolute; for a cell delivered louder it stays at the stamp's own ratio — never extrapolated past what the calibration measured, since the rig's own distortion rises with drive.

drive window The floor and ceiling of the lattice. The hardware window's ceiling is the loudest drive that cannot clip the converter once the input is staged for the pedal, its floor a fixed travel below; the plugin window's ceiling is digital full scale (the float render path has no converter) and its floor is lower, where high-gain preamp knees sit.

edge mask The two columns that sit within a quarter tone of the grid's edges — the sweep's fade-in and fade-out ramps — masked from the tiles and the presence verdict; derived from the grid alone, because the artifact reproduces on a jumpered wire.

excited band The range of fundamentals the sweep drove at full amplitude: from the instantaneous frequency where the fade-in ends to the one where the fade-out begins. Outside it the device was driven under a ramp and a read describes the measurement, not the device.

export grid The summary's frequency columns: a fixed count of fundamentals log-spaced from the sweep's start to its end inclusive, so the first and last columns sit at the sweep's edges, inside its fade ramps.

extraction window (T_k) The raised-cosine window that cuts order k's pulse out of the impulse response; its half-widths are half the gap to each neighboring pulse, capped, and its duration sets the order's spectral resolution.

fine window An optional span of the probe grid filled at a finer spacing (clamped inside the sweep, capped in count so a wide window widens its own spacing), merged into the uniform lattice and sorted, so a transition that falls between two coarse rungs gets several points.

floor class Each point's THD read against the applicable floor: clear (at least the stated margin above it — distortion), noise-dominated (above it but within the margin — distortion mixed with rig noise), or below the floor (at or under it — the read is the rig).

floor runs The THD trace split into maximal runs by floor class, each segment between adjacent points wearing the worse endpoint's class and the boundary points shared, so the solo and compare charts draw one segmentation.

floor source Which of the three bounds the drawn floor came from — the capture's own noise stamp, a null-run reference, or a plugin's stated floor — named only when that source actually sets the floor at some point; the applied floor is always the largest of the available bounds.

floor span The default quietest level, defined as a fixed span (22 dB) below the voltage anchor rather than as a dBFS constant — a constant voltage at the device on every calibrated rig (13.3 mV peak, 12 dB below the Reference Box P6's measured dead-zone handover), −48 dBFS only as its no-volts expression; the ceiling stays an absolute rig-side dBFS.

floor-bound zone The first probe step above the floor: a fitted breakpoint landing in it cannot be anchored from below and is reported as a bound at that level, never as a point.

frame fit The single time shift Δl between the transfer capture's alignment and the paired record's, found by maximizing the magnitude-weighted phase coherence of every fittable harmonic; its residual over the harmonics is the pairing statistic.

fundamental (f_0) The excitation frequency whose harmonics are being read; in a swept measurement every f_0 in the band is excited in turn, and the k-th harmonic of f_0 is read at the output frequency kf_0.

fundamental ellipse The least-squares ellipse through the trajectory: the projection of both cycles onto the fundamental. Its area is the linear-phase loop — what a filter's rotation of the fundamental contributes, never memory.

fundamental presence Whether $H_1(f_0)$ is a measurement at all: absent when it sits at or beyond the analyzer's measured separation bound below the loudest valid in-band order, or at or below its rig-floor margin. Where absent, no ratio to it is quoted.

gain map The level-resolved harmonic measurement (§6.2's gain journey): the Harmonic Distortion sweep repeated at each drive level of a lattice, each level analyzed on its own, and the results summarized on one export grid as thd$[i][j]$ and $H_k[i][j]$ — a map of the device's distortion against level and note.

harmonic order (k) The integer multiple of the fundamental at which a harmonic sits; order 1 is the fundamental itself. Odd orders are produced by symmetric nonlinearities, even orders require asymmetry.

harmonic pulse In the deconvolved response, the impulse response belonging to one harmonic order, arriving $L \ln k$ seconds before the linear one; each is cut out with its own window.

harmonic response ($H_k(f)$) The complex, input-normalized amplitude with which the device emits energy at output frequency f as its k-th harmonic; the harmonic of a fundamental f_0 is read at $f = kf_0$.

harmonic role The tiles' partition of the resolved products by what each does to the requested chord: the recipe (m, n) is placed at index $ma + nb$ on the harmonic grid of the interval's simplest 5-limit just ratio $a : b$, its pitch class read off the odd part of the index; on a played note's class it thickens (at or above the low note) or adds a sub-octave; on a consonant other class it adds color; elsewhere, or between frets, it clashes. Classified at the requested interval, measured at the lattice.

harmonic shaping The read-time predicate deciding whether a harmonic read is content the circuit deterministically shaped or noise that landed in that order's window, from the coherence of the complex response along frequency.

hinge fit The two-segment least-squares fit of output dB against input dB: $y = c_0 + c_1 x + c_2 \max(0, x - k)$, two lines meeting continuously at a breakpoint k swept over a fine set of candidates between the second and the third-from-last probe, each candidate solved by normal equations with a small ridge; the best candidate's slopes are the pre- and post-knee slopes.

impulse response ($h[n]$) The deconvolved response, circular over the analysis length N: the linear pulse at index 0 and order k's pulse at $N - \text{round}(L \ln k \cdot f_s)$.

incremental slope Between consecutive points, the change in output dB per change in input dB: 1 is clean (one dB out per dB in), 0 is a hard limit; the slope arriving at each point.

input-normalized Divided by the sweep amplitude A, so that a unit-gain linear device reads $H_1 = 1$ (0 dB) and every H_k is a ratio to the input, not an absolute level.

instantaneous frequency ($f(t)$) The frequency the sweep is passing through at time t: $f(t) = f_1 e^{t/L}$, the derivative of the phase divided by 2π.

intermodulation product ($mf_1 + nf_2$) A component the nonlinearity places at $mf_1 + nf_2$ for integers m, n; its recipe is the pair (m, n) and its order $|m| + |n|$ is the lowest polynomial power that can produce it. Read at one bin, exactly, and quoted in dBc re the louder played tone.

inverse filter ($\tilde{X}^{-1}(f)$) The closed-form stationary-phase inverse of the synchronized sweep, $2\sqrt{f/L}\,\exp(j(\pi/4 - 2\pi f L(1 - \ln(f/f_1))))$, tapered to zero below f_1; using the closed form extrapolates the phase law above f_2, so harmonics above the sweep's top deconvolve cleanly.

just ratio The simplest ratio $a : b$ with every factor 2, 3 or 5 within 20 cents of the requested 12-TET interval, found by search (3 : 2 for the fifth, 5 : 4 for the major third): the harmonic grid the roles are placed on has fundamental f_1/a, 55 Hz for A2 + E3.

knee The breakpoint of the hinge fit, quoted as a POINT (an input level with a resolution) when the fit is anchored by a unity segment below it, or as a BOUND (at or below a level, with the probe step as its resolution) when the curve is already compressing at the quietest probe or the breakpoint sits within one probe step of the floor.

lattice snap Placing the two tones on bin centers so that a whole number of cycles of each fits the window: the reduced ratio is chosen once against every standard grid (no collisions, spacing at or above the bar, both tones inside a ±15 cent budget, then the interval nearest the request) and scaled to the grid at hand; a grid too coarse for it falls back to the widest collision-free lattice in budget, and a request that cannot escape exact degeneracy is refused before anything plays.

lattice spacing (g) $\gcd(k_1, k_2)$: the integer combinations of the two tone bins are exactly the multiples of g (Bézout), so no product can sit nearer than g bins to another or to either tone. 22 bins on the shipped fifth at 96 and 48 kHz, 24 at 88.2 and 44.1 kHz.

leakage map The parity's own instrument: the pure phantom [1] — a fundamental and nothing else — through the journey's geometry, whose every higher-order read is the analyzer's leakage into that order's window at that column; the allowance every truth comparison subtracts.

level lattice The drive levels a journey measures, evenly spaced in dB from a floor to a ceiling inclusive, one sweep per level; the hardware and plugin windows set the two ends and the level count.

local floor The capture-time noise estimate beside one bin: the median magnitude over the bins from 3 out to about ±8.8 Hz out on either side, with every bin the lattice claims — products, tones, DC, the detector's bins — removed by name. A product is called present when it stands more than four times (12.04 dB) above it; the verdict is stored and never revisited.

measurement floor The level below which a read is the rig rather than the device, derived at read time from the embedded calibration's pooled distortion stamp with per-track, drive and averaging adjustments; a read at or below it is reported as a bound.

median read How the Gain Map reads the calibration stamp: each point replaced by the running median of its neighbors over an odd window before the interpolation, so a single-read trough cannot set the floor; the per-harmonic path reads the stamp verbatim (a window of one).

memory verdict The three-way reading taken on the nonlinear loop area (of the compensated figure when the phase removal is accepted): memoryless-ish below the first threshold, memory effects at and above the second, and a deliberately hedged band between them, where tone filtering ahead of the clipper rotates every harmonic by k times the drive's rotation and a single fundamental cannot separate that from mild memory. Both thresholds are literals in the shipped verdict, which the manifest reads by bisection.

memoryless lattice The set of harmonic phases $\{0, \pi\}$ (in the cycle's convention, after the Novák factor $(k-1)\pi/2$) at which a memoryless device's harmonics sit; each measured harmonic's excess over its nearest lattice point, folded consistently from low order to high, is the filter burden the removal takes out.

methods manifest A machine-readable JSON file, emitted by the shipped measurement kernel through `analysisdump manifest`, listing every parameter that determines a measurement's result — each value obtained by calling the function that owns it. The documents' tables are generated from it and CI fails when the committed copy drifts from a fresh one.

noise bound The noise stamp expressed as a worst-case THD at one point, $\sqrt{2}\,\mathrm{rms}/a_1$: if every bit of the broadband noise power landed in the harmonic bins the THD numerator would read exactly that, so a read at or below it is indistinguishable from the rig; it climbs toward the quiet end as the fundamental shrinks.

noise margin and badge A tile's output rms against its row's noise stamp in dB; below the 6 dB clear margin the tile wears the "≈ noise" badge, whose sentence says most of the drawn wave is the rig's and device's own noise — never that the tile is below a floor, since the wave is drawn.

noise stamp The rms of the capture's own silent pre-roll — the loop plus the device's hiss, measured in the same capture — stored on the curve as its noise floor; the per-step gate and the noise bound both read it, and no calibration stamp enters the compression floor.

nonlinear loop area The unsigned lobe area of the residual trajectory after the fundamental ellipse's output is subtracted — the harmonic content's own lobes, the headline memory figure the verdict is taken on.

note anchor The standard drive note (E2 at the app's 82.4 Hz), guaranteed on every note lattice whose span holds it — the cross-check guarantee — so a matrix always shares a point with the transfer curve's measured cycle at the standard condition; a single canonical note, never per record.

note lattice The matrix's frequency rows: 12-TET-snapped notes over a span, seeded by the endpoints, the standard note where the span holds it and the span's middle note, then refined per side around the middle note so that smaller lattices are subsets of larger ones.

null-run reference A loop-only control capture of the Compression kind under the same calibration, offered as a second floor: its added-harmonic AMPLITUDE (its THD times its fundamental) interpolated between its rungs and divided by this capture's fundamental — an absolute contribution, never its ratio read across — and only inside its own probed span.

one-bin estimator The amplitude reader: a Hann-windowed single-frequency DFT bin with the window's coherent gain compensated, $2\,|\sum xwe^{-j\omega n}|/\sum w$ — exact on the snapped window for a pure tone, and robust to the ramp's tail and to a slightly wrong latency cut.

operative floor A plugin record's own noise bound: the larger of its silence render's rms and its A/A residual, stamped at run initiation; it enters the floor as an absolute dBFS bound with no chain-gain term, and governs where the render-path calibration is numerically silent.

orientation resolution A plugin-path capture is aligned by correlation against a tone and may land half a cycle off; when the cycle's fundamental phase sits on the $\{0, \pi\}$ lattice and the paired polarity confidently disagrees, the output cycle is advanced half a period, restoring the true input–output pairing.

pairing residual The magnitude-weighted RMS phase residual, in radians, over the harmonics of order two and above after the frame fit; a genuine pairing sits well under the accept bar and a record of another session or setting, or a device whose harmonics have no fundamental-anchored phase relationship, does not.

parity test A continuous-integration test that runs the shipped code and the independent Python reimplementation on identical synthetic devices and asserts, at every valid point, agreement between the two to a stated tolerance and agreement of both with analytic ground truth to a stated tolerance.

peak tile The map's headline figure, in three states: the loudest corroborated clear cell as a peak; else the loudest unmasked cell as a bound ("at or below"); else, when the map's fundamental is absent, no percentage at all.

per-step noise gate A probe step whose measurement-window rms sits within the stated margin of the noise stamp is excluded from the curve and counted: its "output level" would be the hiss, not the tone, and would flatten the quiet end into a phantom knee; the quietest surviving step is then the effective floor.

phantom capture A synthetic capture $\sum_k m_k A \sin(k\varphi(t))$ built from the sweep's own phase law, so that its k-th harmonic response is exactly m_k by construction, with no content above the listed orders and no aliasing; the oracle's exact case.

phase removal The § 6.5 correction: the capture's harmonic phases are fitted against the paired Harmonic Distortion record's branch phases with one frame shift, each harmonic's burden over the memoryless lattice is folded out sequentially, and the correction is applied as a full-band unity-magnitude all-pass, so filter phase before and after the clipper is removed and only unpredicted phase — memory evidence — remains.

polarity Whether the device inverts: read from the paired record's H_1 as the energy-weighted mean unit phasor over the low band, whose real part is the confidence (a quadrature response is coherent yet polarity-mute); stated on the figure only when the frame fit confirms the pairing.

presence gate The analyzer's per-bin criterion, amplitude $> 4\times$ the local floor ($20 \log_{10} 4 = 12.04$ dB), decided at capture and frozen in the payload as the product's `aboveFloor` flag.

probe grid The lattice of input levels a compression run steps through: a uniform ladder in dB from a floor to a ceiling, with an optional fine window merged in at a denser spacing; the grid is the user's per run and is read from a record's own points, never assumed.

product order ($|m| + |n|$) The order of the recipe (m, n): the polynomial power that first produces it. The analyzer enumerates every recipe of order two through seven — 54 on a lattice with no collisions — and a shared bin is kept by its lowest order.

product reading One of four outcomes per product, decided in a fixed order: refused on lattice position (beside a played tone or DC), out of band (under 20 Hz), refused beside a louder product, refused by the analyzer's window or under the loudest product's skirt, then present with its margin over the capture floor or absent with the bin's own read as a bound. A refusal is never silently absent.

rate constant (L) The time in seconds over which the sweep's instantaneous frequency grows by a factor of e: $f(t) = f_1 e^{t/L}$. Fixed by the synchronization condition from the requested duration.

read scatter The one-sigma uncertainty, in dB, of a magnitude read at its estimated signal-to-noise ratio — the Rician mean-of-dB statistic, pinned by a Monte-Carlo table and saturating at the pure-noise value.

reduced tone ratio ($a : b$) The tone bins in lowest terms. Two recipes of order at most K can share a bin only when $a + b \leq 2K$, so the probe insists on $a + b \geq 2K + 1$; the ratio alone fixes the delivered interval, whatever the rate or length. The shipped fifth is $109 : 163$.

refinement (nesting) Building a lattice by inserting points into the widest gap of an existing lattice rather than dividing the span evenly, so that every smaller lattice is a subset of every larger one ($3 \subset 5 \subset 7$) and matrices of different dimensions compare point for point; an even divider misses itself across dimensions.

residual alignment error Any difference between the integer loop latency removed before binning and the true delay, in samples; it rotates the output's fundamental by $2\pi f_0 e/f_s$ and harmonic k by k times that, opening the figure into a phase loop that grows with the note.

resolution band For a point knee, the interval of candidate breakpoints whose fit error stays within five percent of the best; half its width is the knee's resolution, never finer than the candidate spacing, and it sets the decimals the level is displayed with.

row noise stamp Per frequency row, the rms of the row capture's silent pre-roll (the loop plus the device's own hiss), stored beside the tiles as the row's noise bound; computed inline at both capture sites by one expression (a mirror the manifest states) and never used to gate the stored samples.

sample array (encoding) A tile's stored channel: Float32 little-endian bytes (base64 in the payload JSON) with the encoding name, the channel and the sample count stated explicitly; a reader refuses an unknown encoding or a byte count that disagrees with the stated count rather than guessing a layout.

settle region and measurement window Each step's two parts: the settle region, in which the amplitude ramps to the step's value and the device's response is allowed to settle (not read), and the measurement window that follows, over which the estimator reads the tone.

sideband pair statistic The mean of the two innermost detector bins about each tone, and the difference between the two tones' pairs against $20\log_{10}(f_2/f_1)$: a shared fractional frequency wander separates the pairs by that figure, while amplitude or frequency-independent phase modulation leaves them equal — a shape the line states without a verdict.

static curve The output-versus-input estimate that ignores the loop: the cycle is split at the input's extremes into a rising and a falling branch, each is interpolated onto a common input grid, and the two are averaged — never binned by input value.

stepped tone The stimulus: a single sine at the probe note whose amplitude steps through the grid, each step a settle region (with a raised-cosine ramp from the previous amplitude) followed by a measurement window, the phase continuous across steps, a final fade to zero over one ramp.

summed presence reading The fundamental-presence reading at one compression point, with the SUMMED harmonic content as the reference in place of the loudest order (the payload stores only the sum) — at most $10\log_{10}(K-1)$ dB more eager than the per-order form, conservative in the refuse-to-quote direction; the majority over the points gates the knee, the cleanup and every quoted THD.

sweep averaging Combining N back-to-back sweeps from one capture by complex, bin-wise averaging of their harmonic responses after each repeat is derotated by its fitted capture delay; the noise power falls as $1/N$ where the residual is noise, and magnitudes are never averaged.

synchronization condition The requirement that $f_1 L$ be an integer, met by rounding the rate constant; it makes the sweep's phase at every harmonic's arrival a whole number of cycles, so harmonic responses separate with meaningful phase.

synchronized swept sine ($x(t)$) A sine whose frequency rises exponentially from f_1 to f_2 over the duration T, with its rate constant chosen so that $f_1 L$ is an integer; under that condition each harmonic a memoryless device generates is an exact time-advanced copy of the sweep, advanced by $L \ln k$.

tanh The hyperbolic tangent, as the saturating function $y = \tanh(gx)$: linear for small inputs and bending smoothly toward ± 1, which it never reaches — the textbook soft clipper. It is odd-symmetric, so it produces odd harmonics only, and its ladder falls off steeply with order. The worked example throughout these documents is $\tanh(2x)$ at a sweep amplitude of 0.5, whose exact harmonic amplitudes are computed by quadrature.

tile One cell of the matrix: two sample arrays, the stimulus and the capture over the same time span — a whole number of periods of the row's note — starting at an upward zero crossing of the stimulus, the output at the loop's latency so the device's phase shift is preserved.

total harmonic distortion (THD) $\sqrt{\sum_{k\geq 2} |H_k(kf_0)|^2}/|H_1(f_0)|$ over the orders that are in the audible band and pass the validity bound — THD, not THD+N; undefined where the fundamental is not present.

transfer curve The X–Y figure of a device under a steady low-frequency sine: instantaneous output against instantaneous input over one phase-averaged cycle. A memoryless device retraces a single curve; linear filter phase and genuine memory open it into a loop.

truncated column A column where the validity bound removed an order that the declared band would have counted, so its THD cells under-read their band; the column is flagged and excluded from the tiles, never corrected.

two-tone measurement Chord IMD: two simultaneous sines at a musical interval, snapped to exact bins of the analysis window, driven through the device; the result is the level of every intermodulation product up to the seventh order, the three capture-time percentage sums, and the coherence detector's reads beside each tone.

unity guard The pre-knee test: a fitted lower segment more than a stated tolerance under 1 dB/dB means no clean segment anchors the fit — the device is already compressing at the quietest probe — so the knee is a bound at the floor with the fitted slope quoted as evidence; the test is against unity, not against the compression threshold.

unsigned lobe area The total loop area: the closed trajectory is split at every self-crossing and the absolute area of every lobe is summed, so opposite-winding lobes add instead of canceling; lobes under a floor fraction of the box are dropped as bin noise.

validity bound Where a harmonic read is trustworthy: $f_1 \leq f_0 \leq f_2$ and $f_0 \leq f_s/(2k+1)$ (the alias-image bound), plus $f_0 \leq f_{\text{cal}}/k$ on a compensated record. Reads outside it are excluded from every fit, feature, summary and comparison.

value alignment (app level) The Compare view's rule for two or more matrices: tiles are matched on the STORED delivered notes and levels (1.7 cents, 0.05 dB), never on row or column position, the level basis decided by the same predicate that draws the axes — dBFS on a shared axis, delivered voltage across calibrated rigs, no level join across domains — and the alignment always stated.

voltage anchor The standard drive as a VOLTAGE at the pedal's input (168 mV peak), expressed in a rig's own dBFS through its measured volts factor ($20 \log_{10}(0.168/V_{\text{FS}})$) or as -26 dBFS by convention where no factor is known; the level lattice's middle point and the reference its default floor hangs from.

waveform matrix The measured waveform matrix: raw captured input and output tiles over a note × level grid, stored as the capture itself (Float32 at the capture rate, a fixed number of periods each, no resampling). It has no estimator; the method is the choice of which samples are kept and how they are labeled.

zero-crossing start The tile's first sample: the first stimulus sample at or after the step's measurement window whose phase is a whole number of cycles (the input rising through zero), so every tile begins on the same phase of the input; a start-index choice, never a resample, and the output tile is the same span at the loop's latency.

A Scripts, tests and how to reproduce

Paths are relative to `Docs/Methods/` in the repository, except the tool, which lives under `Tools/`. The repository: https://github.com/PhaseDog/PedalScope

Path	Role
`Tools/analysisdump`	the shipped-code oracle: `synth` (a known device through the shipped analyzer, with exact truth beside it), `transfer-synth` (a known device through the shipped transfer path, paired, with its truth beside it), `compression-synth` (a known device through a known loop and the shipped Compression path inside the app's assembly, the ladder beside every step), `journey-synth` (a known device through a known loop and the shipped Gain Map path, the ladder beside every cell), `imd-synth` (a known device or a phantom through the shipped two-tone path, with the torus-quadrature truth beside every product), `matrix-synth` (a known device through the shipped Waveform Matrix plan end to end — lattices, stimulus, cut, payload — with the exact truth beside every tile and device state carried across rows), `windows`, `manifest`
`generated/manifest.json`	the manifest, committed, drift-checked
`py/harmonic_distortion.py`	the reimplementation; `-help` lists the same options as `synth`
`py/parity_hd.py`	the comparison and the tolerances, with their provenance
`py/test_parity_hd.py`	the CI parity test (pytest)
`py/transfer_curve.py`	the Transfer Curve reimplementation (the options of `transfer-synth`)
`py/parity_transfer.py`	its comparison and tolerances, with provenance
`py/test_parity_transfer.py`	its CI parity test
`py/compression_curve.py`	the Compression reimplementation (the options of `compression-synth`); imports `harmonic_distortion.py` and `transfer_curve.py`
`py/parity_compression.py`	its comparison and tolerances, with provenance
`py/test_parity_compression.py`	its CI parity test
`py/gain_map.py`	the Gain Map reimplementation (the options of `journey-synth`); imports `harmonic_distortion.py`
`py/parity_journey.py`	its comparison and tolerances, with provenance
`py/test_parity_journey.py`	its CI parity test
`py/chord_imd.py`	the Chord IMD reimplementation (the options of `imd-synth`)
`py/parity_imd.py`	its comparison and tolerances, with provenance
`py/test_parity_imd.py`	its CI parity test
`py/waveform_matrix.py`	the Waveform Matrix reimplementation (the options of `matrix-synth`); imports `harmonic_distortion.py`, `transfer_curve.py` and `compression_curve.py`
`py/parity_matrix.py`	its comparison and tolerances, with provenance
`py/test_parity_matrix.py`	its CI parity test
`py/make_figures.py`	the figures, seeded, byte-stable, with TSV sidecars of the plotted values
`py/gen_docs.py`	the generated fragments this document inputs
`terms.yaml`	the terms source; `py/test_terms.py` is its guard

Path	Role
`tex/changes-methods.tex`, `tex/changes-explained.tex`, `tex/release.tex`	the two changes appendices and the title-page stamp; `py/test_changes.py` is their guard
`make methods-test`	parity and terms guards
`make methods-docs`	manifest → figures → fragments → both PDFs

Toolchain: Python 3 with the versions pinned in `Docs/Methods/requirements.txt`; LaTeX through `latexmk` with `pdflatex`, pinned in `Docs/Methods/tex/latexmkrc`. The Swift tool builds with the repository's own toolchain.

Changes since the previous version

Each entry lists what changed in this document against the copy published before it, newest first. The plain-language companion keeps its own list, entry for entry.

1.1 (43) — 30 September 2026

Against the copy published on 23 September 2026, the first public copy.

- **Harmonic Distortion.** The top of the excited band is now part of every order's validity bound on every analysis — raw, simulated and compensated alike — so the grid point on f_2, inside the fade-out, is outside the band for every order (before, only a compensated record's calibrated-band bound removed it). The THD curve delivers all 128 requested points, to 10 kHz; it delivered 127, ending at 9552.89 Hz.
- **Synthetic devices.** Every synthetic device in the parity runs is applied at four times f_s through the anti-alias chain the app's Simulation Mode uses (`synthesis.oversampling` in the manifest), replacing application sample by sample at f_s. The per-device tolerances and the band-top limitation are restated for that chain: its roll-off, not folded content, is what the tops of H_7's and H_9's validity bands read. The reimplementation gains a full-wave rectifier source ($|x|$), matching the kernel's.
- **Transfer Curve.** E6 joins the probe notes, between A5 and A6; the table lists eleven.
- **Chord IMD.** An over-bar coherence reading no longer declares every product untrustworthy. The new paragraph "What the verdict means for each product" (Section 7.4) predicts the smear at each product's bin from the outer detector reads and from the loud products' own skirts, reads a product that stands the smear margin clear of it, and marks the rest unread — neither present nor absent, their tiles lower bounds, their quiet bins out of the floor. The parity table and the limitations gain the model's terms; the legacy stored-products path and an uncleared lattice keep the earlier sentence.
- **Waveform Matrix.** The 7×7 lattice's tie is broken toward the quiet side: two gaps whose widths differ by less than the manifest's `gapTieToleranceDB` are the same width. The section records the rounding residue that decided the tie before, the two factors that delivered different lattices, that 3×3 and 5×5 lattices are unchanged, and that a stored lattice never moves.
- **Every chapter.** "Purpose, and what the measurement cannot tell you" is now "Significance and limitations"; the Transfer Curve's and the Waveform Matrix's state that every measurement here reads the device at a constant stimulus amplitude after settling. The first section is "How to read this document". Spelling is American throughout, several glossary entries are reworded, and the glossary gains "bound".

References

[1] Angelo Farina. Simultaneous measurement of impulse response and distortion with a swept-sine technique. In *Audio Engineering Society Convention 108*, Paris, February 2000. Preprint 5093.

[2] Antonin Novák, Laurent Simon, František Kadlec, and Pierrick Lotton. Nonlinear system identification using exponential swept-sine signal. *IEEE Transactions on Instrumentation and Measurement*, 59(8):2220–2229, 2010. doi:10.1109/TIM.2009.2031836.

[3] Antonin Novák, Pierrick Lotton, and Laurent Simon. Synchronized swept-sine: Theory, application, and implementation. *Journal of the Audio Engineering Society*, 63(10):786–798, 2015. doi:10.17743/jaes.2015.0071.

[4] S. O. Rice. Mathematical analysis of random noise. *Bell System Technical Journal*, 23(3): 282–332, 1944.