


TECHNICAL NOTE



THIRDEAR


AUDITORY DISTORTION PRODUCT INSTRUMENT

Model of 2026

An instrument whose output register is empty · the nonlinearity that produces a tone absent from the stimulus · how the phase of the stimulus sets the tone's level · one equation for spontaneous emissions, compression and distortion · and why two distortion products need two different mechanisms

DATABURN

AUDIO RESEARCH
DATABURN.ORG



PSYCHOACOUSTIC SYNTHESIS

This instrument never outputs the note you play. What leaves the plugin is a comb of six partials clustered around a carrier, spaced by exactly the frequency of the note you asked for. The note itself is not transmitted, and it cannot be: any energy at that frequency would put the note in the signal as a plain sine wave, and there would be no illusion.

The listener hears it anyway, because the ear is not a linear device. A compressive, asymmetric nonlinearity applied to a sum of tones produces sum and difference frequencies that were not in the input. Model that with the first terms of a power series, $y = x + a_2x^2 - a_3x^3$. Because the comb is exactly harmonic, the even term's difference tone for every adjacent pair lands on the same frequency, and they sum there. The odd term's products fall at $(m + i + j - k) \cdot fp$ for partial indices up to $N - 1$, so for a comb whose lowest harmonic number m exceeds the partial count N (fp below carrier/ $(N + 0.5)$; here $m = 27$ against six partials) it cannot reach fp at all, and removing it leaves the phantom unchanged. Above that boundary it can: at C5 the comb is harmonics 6 to 11, and at these coefficients the cubic contribution sits 13 to 42 dB under the quadratic one across the instrument's range of partial counts and harmonic numbers. The gap is smallest with two or three partials on the lowest harmonics: for a pair at harmonics 2 and 3, the CDT $2f_1 - f_2$ and the QDT $f_2 - f_1$ are the same frequency. Removing the even term erases the phantom everywhere. Page 7 runs the same test on a model that has only the cubic term.

■ A 110 Hz PHANTOM FROM A 3 kHz COMB




fig. 1: The same second of audio, before and after a memoryless compressive nonlinearity. The phantom register is empty to the numerical floor on the left and carries a component on the right. Nothing was added to the signal: the nonlinearity is the only difference between the two panels.

	LEVEL AT fp	READING
emitted signal	-180 dB	the numerical floor: there is no energy at the phantom frequency, by construction.
after the ear	-13 dB	generated inside the listener, from partials more than four octaves above it.

PHASE

The distortion product is not generated once. Every adjacent pair in the comb contributes its own difference tone at the same frequency, and what the listener hears is the vector sum of all of them. A vector sum depends on phase. Pressnitzer and Patterson measured this directly: with the primaries in cosine phase the contributions add, and each doubling of the number of pairs adds about 3 dB. Scramble the phases and they partially cancel. This is the reason the instrument resets every oscillator phase together at note-on, and the reason Phase is a control on the panel. Two hundred random-phase trials against the phase-locked stimulus put the cost of scrambling at 8 dB of phantom level.

WHAT PHASE SCRAMBLING COSTS THE PHANTOM




fig. 2: Two hundred trials with the six partials given uniformly random phases, against the phase-locked stimulus the plugin actually emits. Scrambling never helps and typically costs about 8 dB. The distribution has a long left tail: some phase draws cancel the phantom almost entirely.

PHASE SETTING	WHAT IT DOES	MEASURED
Locked	all partials in cosine phase; contributions add	-41.2 dB, the reference
Alternating	every other partial shifted by $\pi/2$. Pressnitzer and Patterson's experiment 3: the odd components of the distortion spectrum cancel, leaving a spectrum built on 2fp, so the percept rises an octave. With six partials, one of the five pairs has no partner to cancel.	-55.2 dB, 14 dB down (five pairs: 14 dB in closed form)
Scrambled	uniformly random phases per note	-49.2 dB median, 8 dB down

LIMITS OF THIS MODEL

The page 2 nonlinearity is memoryless, so every pair of primaries contributes its difference tone with the phase difference of that pair and nothing else. It gets the alternating-phase experiment right for exactly that reason: a quarter turn on every other partial makes adjacent pairs differ by plus and minus a quarter turn, and their difference tones cancel in twos. (A half turn does nothing, in this model or any other: the comb with every other partial flipped is the locked comb delayed by half a period of fp, and the first release of this instrument shipped that by mistake.) What the model gets wrong is the buildup. Doubling the pair count adds a full 6.0 dB here against about 3 dB reported in listeners. One plausible reason is that the travelling wave reaches each place at a different time, so a real cochlea gives each pair its own phase and the contributions only partly reinforce; Pressnitzer and Patterson also report the product's phase nearly constant across partial counts, so the mechanism is not settled. Whichever it is, it is an ingredient this model lacks.

THE HOPF OSCILLATOR

The power series on page 2 accounts for the phantom and nothing else. It cannot explain why the ear amplifies quiet sounds far more than loud ones, and it cannot explain why a healthy ear emits sound of its own. Those are usually presented as three separate properties of the cochlear amplifier: active, nonlinear, sharply tuned. One equation produces all three.

$$\dot{z} = (\mu + i\omega)z - |z|^2z + F(t)$$

the normal form of a supercritical Hopf bifurcation [5]

z is the displacement of one section of the cochlear partition, ω its natural frequency, and μ the control parameter that decides whether the section is damped or self-oscillating. Below zero it is silent and needs a stimulus. Above zero it oscillates with no input at all, which is a spontaneous otoacoustic emission: a sound the ear makes rather than receives. Most normal adult ears have one, and women more often than men. The cited study found 83 per cent in women and 62 in men once it had looked every way it could, and anywhere from two-fifths to four-fifths depending on how it looked [7]. The square-root approach to threshold is the signature of the supercritical case, where the amplitude grows continuously from zero rather than jumping.

■ ONSET OF A SPONTANEOUS EMISSION




fig. 3: Sixty oscillators integrated from rest with no input, one per value of the control parameter. Below threshold every one decays to silence; above it they settle onto a limit cycle whose amplitude follows the square root of the distance past threshold. The simulated points sit below the theory just past threshold, and that gap is not a plotting artefact: the time constant there diverges as one over the control parameter, so a fixed-length run has not finished growing. It is the same critical slowing down measured on the next page, showing up here as a visible bias.

COMPRESSION

Hold the same oscillator exactly at the critical point and drive it. On resonance the fast rotation drops out and what is left is the amplitude equation $dr/dt = \mu r - r^3 + F$. A passive resonator answers proportionally: double the drive, double the response. The critical oscillator answers as the cube root of the drive [8]. A cube-root law maps 120 dB of acoustic range into 40 dB of response, which is roughly what a hair cell can deliver, and it is why quiet sounds receive enormous gain and loud ones almost none. The power law is the settled part; its biological implementation is still under investigation.

■ RESPONSE AGAINST DRIVE, LOG-LOG




fig. 4: Steady state solved exactly rather than integrated, for the reason in the caution below. Fitted log-log slopes: 0.333 at the critical point against a theoretical one third, and 0.987 for the damped oscillator against a theoretical one. At the quietest drive the critical oscillator answers 61 dB louder than the passive one.

CAUTION

Settling time at the critical point diverges as $F^{-2/3}$, which is critical slowing down. At the quietest drive on this plot the time constant is 718 model-seconds against 0.72 at the loudest, a factor of 1000. A fixed-length integration silently under-reports the quiet end; the first attempt at this figure measured a slope of 0.69 for exactly that reason. The slowness is not a numerical nuisance either. It follows from operating at the critical point, and it is one reason faint sounds take a moment to become audible.

SPONTANEOUS EMISSIONS

A real cochlea is not one oscillator but a graded array of them, and no two are identical. That irregularity is called cochlear roughness and it is what makes every ear measurably unique. Here are sixty oscillators with smoothly graded frequencies, small random scatter in both frequency and control parameter, weak coupling to their neighbours, and no input whatsoever.

■ SIXTY OSCILLATORS, NO INPUT




fig. 5: Top: the oscillators that crossed into self-oscillation, in accent. Bottom: the sound in the ear canal, which is the sum of everything. It holds a handful of discrete, stable tones, irregularly spaced but never too close together, which is the observed character of real spontaneous emissions. The floor on that spacing is the coupling: rerunning the identical bank with the coupling removed gives 11 peaks with a minimum spacing of 0.075, against 7 peaks at 0.150 here. Neighbours that fall too close together entrain and emit as one.

What this establishes is narrow: a heterogeneous coupled-oscillator field with no input produces a small number of discrete, stable, irregularly spaced emissions, and the coupling sets the minimum spacing. It does NOT establish the mechanism behind measured emission spectra, which ripple by as much as 20 dB peak to peak and are attributed to interference between two cochlear sources rather than to entrainment between neighbours [4]. Both trace back to mechanical irregularity; only one of them is modelled here. For the player the consequence is the same either way: the phantom is not equally strong at every pitch, and where it is strong differs between listeners. Each listener hears the instrument through the irregularities of their own cochlea.

TWO MECHANISMS

Look again at the nonlinear term in the normal form. $|z|^2 z$ is purely cubic. Drive a critical oscillator with two tones and it must produce intermodulation at $2f_1 - f_2$, the cubic difference tone, which is the strongest auditory distortion product and the one this instrument's CDT mode places deliberately. So the model that explains why some ears emit spontaneous tones also explains why every ear hears a cubic difference tone. Both come from the same term.

TWO TONES IN, DISTORTION PRODUCTS OUT




fig. 6: A single near-critical oscillator ($\mu = -0.02$) whose characteristic frequency is the CDT place, driven by two tones at the clinical ratio of 1.22. The cubic difference tone arrives at -6.3 dB. The quadratic difference tone does not arrive at all: it sits at -91 dB, in the numerical floor. The level is not a prediction: it is set by the drive amplitude, the detuning and the oscillator answering at its own place. Goldstein's -15 dB re primaries, which the panel draws, is the listener's figure; the result here is the parity, not the number.

That absence is the most useful result in this document, and it is not a numerical failure. The normal form's only nonlinearity is cubic, so it can generate odd-order intermodulation and nothing else. Real ears produce strong quadratic difference tones, which means they carry an asymmetric nonlinearity this model does not have: the asymmetry of the hair cell's own transduction curve, which is exactly what the $a_2 x^2$ term on page 2 was standing in for.

WHY THERE ARE TWO MODES

The two panel modes are not one effect voiced two ways. QDT Complex relies on transducer asymmetry, which needs many partials and rewards phase locking. CDT Pair relies on the critical oscillator, which needs only two tones and is nearly independent of level. They have different growth laws, different ratio sensitivities and different audibility floors, and for a pair, where the parity argument is exact, neither the quadratic term nor the cubic one produces both. The ear has both mechanisms; the Mode switch chooses which one to drive.

THE RATIO WINDOW

In CDT mode the phantom sits at $2f_1 - f_2$, so the interval between the two primaries decides the pitch of the tone the listener's ear generates, and the three-note chord that results. Kendall, Haworth and Cadiz tabulated the ratios at which that chord is recognisable. They are simple integer relationships, and the Ratio control carries a soft detent at each one.

■ WHERE THE RATIO CONTROL IS USEFUL




fig. 7: The whole useful travel of the Ratio control, on one axis. Below 1.14 the two primaries are close enough to beat against each other and the roughness competes with the phantom; above 1.25 the distortion product drops sharply. The detents are the integer chords; the default sits at the ratio that produces the largest measured emission.

f_2/f_1	CDT : f_1 : f_2	CHORD	NOTE
1.111	8:9:10	whole-tone cluster	below the roughness threshold
1.125	7:8:9	~whole-tone cluster	below the roughness threshold
1.1428	6:7:8	non-tertian triad	
1.166	5:6:7	~diminished triad	
1.2	4:5:6	major triad (m3)	inside the clearly-audible window
1.25	3:4:5	major triad (M3)	top of the window; drops sharply above

CONTROLS

ThirdEar has eighteen parameters. The first fifteen are terms in the models above or documented conditions for the illusion to work; the last three measure it, by the method Goldstein used.

CONTROL	RANGE	WHAT IT SETS
Mode	QDT complex / CDT pair	Which mechanism is being provoked. Complex builds a harmonic comb and relies on transducer asymmetry; Pair emits two tones and relies on the critical oscillator. See page 7.
Carrier	800 Hz - 8 kHz	Where the comb is placed. The phantom's frequency does not change with it; only the register that carries it does.
Partials	2 - 8	How many primaries. Doubling the number of adjacent pairs adds about 3 dB in listeners, but a full 6 dB in the memoryless model of page 2. That gap is the one place the model is known to over-predict; see page 3.
Ratio	1.05 - 1.30	CDT mode only: f_2/f_1 , and therefore the phantom's pitch and the chord. Soft detents at the six integer ratios on page 8.
Rove rate	0.02 - 2 Hz	How fast the carrier wanders. To rove a stimulus is to vary it on purpose, the psychoacoustic term for taking a cue away. The comb re-forms on the nearest whole multiple of the note, so the tones move and the phantom stays where it is: evidence that the pitch is made in the ear. The illusion also fatigues on a static stimulus, and moving it keeps the mechanism from adapting. QDT mode only.
Rove depth	0 - 1 oct	How far the carrier wanders, either side of where Carrier sets it.
Attack	0.5 - 100 ms	Onset time. Measured across the full range it does not change the phantom's level at all, so it shapes articulation and nothing else.
Release	2 - 500 ms	Decay time.
Level	-60 - -6 dB, default -18	Output level. The quadratic product grows about 2 dB per dB of primaries, so relative to them it gains 1 dB for every dB above the 60 dB SPL at which the panel's figure is quoted. The cubic difference tone grows about 1 dB per dB in listeners (Goldstein), where a memoryless cube would give 3. Level changes the balance between the two modes as well as the volume. The ceiling is -6 dB and the synth is capped so its coherent pulse never reaches the safety clip (a clipped pulse puts the phantom in the signal). Work at the lowest level at which the phantom is audible; see page 11.
Norm	peak / energy	Partial normalisation. Peak ($1/N$) holds the summed peak constant; each partial gets quieter as you add more and the phantom falls about 7 dB from two partials to eight in the memoryless model. Energy ($1/\sqrt{N}$) lets it rise about 5 dB over the same span. Neither holds the per-partial level constant.
Pattern	12 patterns	The order the internal pattern engine steps through held notes; the panel labels this list ORDER. A time-varying phantom layer is one of the documented conditions for the effect; 5 to 15 events per second is the productive range.
Rate	2 - 20 /s	Pattern steps per second.
Phase	locked / alternating / scrambled	Note-on phase relationship. Locked is the reference; alternating (a quarter turn on every other partial) cancels the odd components of the distortion spectrum and raises the percept an octave; scrambled costs about 8 dB in the model. See page 3.
Mode seq	11 interleavings	Alternates the two modes across successive notes, so a single line can provoke both mechanisms in turn.
Arp clock	free, 1/1 - 1/32 (10 settings)	What clocks the pattern. Free reads Rate as Hz; a division locks the steps to host tempo. At 120 bpm 1/8 and 1/16 land at 4 and 8 events per second, inside the productive window.
Cancel	off / on	A sine at the phantom's own frequency, phase-locked to the comb: Goldstein's cancellation method [1]. Turn it on, then turn Cancel level and Cancel phase until the phantom disappears. Click the phantom card to open the panel for it.
Cancel level	-60 - -5 dB	The tone re one primary. Where the phantom disappears, this is the product's level in your ear, beside the model's figure on the card. Headphones colour it: the tone and the comb sit in different parts of their response. The product's phase can move with level, so null again after changing Level.
Cancel phase	0 - 360 deg	The tone re a sine at fp that starts with the note. Every adjacent pair of the comb differs by exactly fp, so the product's phase is locked to the tone and the null holds for the whole note and while the carrier roves (tests/cancel.cpp measures both).

Maryanne Amacher (Kane, Pennsylvania, 1938 - Rhinebeck, New York, 2009) studied with George Rochberg and Karlheinz Stockhausen, did graduate work in acoustics and computer science at Illinois, and spent the rest of her life building music out of the room, the building and the listener rather than out of instruments. City Links (from 1967) piped live sound between distant places; Music for Sound-Joined Rooms (from 1980) staged a piece across the architecture of a house; the Mini Sound Series (from 1985) serialised that idea like television. She made the storm for John Cage's Lecture on the Weather and Close Up for his Empty Words, and scores for Merce Cunningham, Torse among them.


From the 1970s on she composed with what she called ear tones: the distortion products a healthy cochlea generates when it is driven by two or more tones at once. She was the first to use them systematically as musical material, and she was explicit that the listener was part of the instrument. The 1999 Tzadik record carries the method in its subtitle.

Her note on the sleeve is the instruction this instrument is built to follow; page 11 puts numbers on how high "quite high" may go:

"When played at the right sound level, which is quite high and exciting, the tones in this music will cause your ears to act as neurophonic instruments that emit sounds that will seem to be issuing directly from your head."



MARYANNE AMACHER AT THE DESK, 1970s
PHOTOGRAPHER NOT YET CREDITED



SHE HEARS A RHAPSODY
CITY-LINKS, WBFO BUFFALO, 26-27 MAY 1967:
28 HOURS ON FIVE TELEPHONE LINES.
CAMPUS PAPER CLIPPING, CREDIT NOT YET ESTABLISHED

SOUND CHARACTERS

(MAKING THE THIRD EAR)

Maryanne Amacher. Tzadik TZ 7043, 1999. Opens with "Head Rhythm 1" and "Plaything 2".
Followed by Sound Characters 2 (Making Sonic Spaces), Tzadik, 2008.

What ThirdEar adds is fifteen years of the literature that followed her. Kendall, Haworth and Cadiz set out in 2014 to reproduce her method with a synthesis engine and a listener model, and their ratio table is on the panel as detents. Pressnitzer and Patterson's vector-sum account is why every oscillator here starts in phase, and their alternating-phase control is a switch on the front. The honesty meter exists because her claim, that the melody is in the listener and not in the signal, is a claim a spectrum analyser can check.

LISTENING

WORK	WHERE TO START	WHY
Sound Characters (1999)	track 1, headphones, the level she asks for and no more	the interlocking high patterns; the low line under them is produced in the ear
Sound Characters 2 (2008)	any track	the same method in rooms rather than heads
Kendall, Haworth, Cadiz (2014)	the audio examples that accompany the paper	the CDT triads the Kendall presets reproduce
Estelle Schorpp, In My Ears (for Maryanne), LINE_155, 2024	synaptic choral I; SuperCollider, pairs of close pure tones, 1 to 4 kHz	the other mode: a close pair puts 2f1-f2 a third of an octave under f1, so the phantom stays high where Amacher's wide combs put f2-f1 an octave or more down

EXPOSURE SOURCES

This instrument works by driving a nonlinearity in the listener. It emits narrowband tones at moderate to high level, held for as long as a note is held, and the effect is strongest exactly where hearing is most sensitive. The effect is an illusion, but the stimulus is not: it is ordinary acoustic energy delivered to a listener's ear, and the ordinary exposure limits apply to it. They are listed below with their sources.

LIMIT	VALUE	SOURCE
Occupational exposure	85 dBA, 8 h time-weighted average. The 3 dB exchange rate means every additional 3 dB halves the allowable duration.	NIOSH 1998, pub. 98-126
Recreational allowance	80 dBA over 40 h per week for adults; 75 dBA for children.	WHO-ITU H.870 mode 1
Rest breaks	Narrowband tones concentrate energy on a small region of the cochlear partition, so take them. No published standard sets an interval for this; the limits above are dose-based.	lab practice, not a standard
A-weighting	Across the full 800 Hz to 8 kHz carrier range the A-weighting correction stays between about -1.1 and +1.3 dB, so dB SPL and dBA are interchangeable here.	IEC 61672 weighting curve

WARNING

Work at the lowest level at which the phantom is audible. Raising the primaries raises the quadratic product about twice as fast in dB, so the illusion gets stronger long before the stimulus gets loud. A phantom that fades during a long session is the mechanism adapting to a steady stimulus: raise Rove depth (QDT mode), or stop for a few minutes. Do not turn the level up to bring it back.

SOURCES

- [1] Goldstein, J. L. "Auditory nonlinearity." Journal of the Acoustical Society of America 41, 1967. Growth laws and audibility floors for both products; the ratio-sensitivity difference that separates the two modes.
- [2] Pressnitzer, D. and Patterson, R. D. "Distortion products and the perceived pitch of harmonic complex tones." Physiological and Psychophysical Bases of Auditory Function, 2001. The vector-sum account, the 3 dB per pair doubling, and the alternating-phase experiment behind the Phase control.
- [3] Kendall, G., Haworth, C. and Cadiz, R. "Sound synthesis with auditory distortion products." Computer Music Journal 38(4), 2014, pp. 5-23. Table 1 is the source of the six ratio detents.
- [4] Mauermann, M., Uppenkamp, S., van Hengel, P. W. J. and Kollmeier, B. "Evidence for the distortion product frequency place as a source of DPOAE fine structure in humans. I." Journal of the Acoustical Society of America 106(6), 1999, 3473-3483. Source of the 20 dB ripple and the 1.22 ratio. Note that this paper argues a TWO-SOURCE origin for fine structure, which is not the mechanism modelled on page 6.
- [5] Hudspeth, A. J., Julicher, F. and Martin, P. "A critique of the critical cochlea: Hopf, a bifurcation, is better than none." Journal of Neurophysiology 104(3), 2010, 1219-1229. The normal form used on pages 4 through 7.
- [6] Kemp, D. T. "Stimulated acoustic emissions from within the human auditory system." JASA 64, 1978. The original observation that a healthy ear emits sound.
- [7] Penner, M. J. and Zhang, T. "Prevalence of spontaneous otoacoustic emissions in adults revisited." Hearing Research 103, 1997, 28-34. The prevalence figure on page 4.
- [8] Eguiluz, V. M., Ospeck, M., Choe, Y., Hudspeth, A. J. and Magnasco, M. O. "Essential nonlinearities in hearing." Physical Review Letters 84, 2000, 5232. The one-third power law and the cubic difference tone from a critical oscillator.

COMPANION MATERIAL

Every figure is recomputed by docs/build_thirdear_note.py from the models in earlab/notebooks/01_phantom_in_silico.ipynb and 03_hopf_oscillators.ipynb, which carry the full derivations and the limits of each model. Measured constants and their citations are in earlab/psychophysics/constants.csv.