ZERO-DELAY FEEDBACK FOR FM CROSS-MODULATION

Luc DÖBEREINER[1] and **Dario SANFILIPPO**[1]

[1]**Freelance Composer**, Vienna, Austria

ABSTRACT

FM cross-modulation — the mutual, bidirectional phase modulation of two coupled oscillators — introduces an implicit algebraic loop that digital implementations conventionally resolve by inserting one-sample delays. We show that this resolution is regime-dependent: it approximates the implicit system where coupling is weak, but in the strongly coupled regime it constitutes a structurally distinct dynamical system with its own bifurcation structure and a persistent $f_s/4$ spectral artefact. Reformulating cross-modulation as a per-sample implicit system — in the spirit of zero-delay feedback (ZDF) methods from virtual analog filter design — we compare four realisations: the one-sample delay, fixed-point iteration (FPI), a trapezoidal-style implicit averaged scheme solved by Newton-Raphson, and a root-tracking solver, the first three evaluated at native sample rate and under $64\times$ oversampling. A per-sample solution always exists; it is guaranteed to be unique for $\beta < 1$, and beyond this boundary multiple solutions may coexist, so that each realisation imposes a branch-selection rule — a musical decision. FPI is globally convergent for $\beta < 1$; the averaged system remains contractive, with a unique solution, up to $\beta < 2$, and its safeguarded Newton solver converged to machine precision throughout the tested range; the root-tracking solver solves the original equations directly and runs in real time. All four are implemented as a SuperCollider UGen and a JUCE audio plugin; source code and audio examples are available in the accompanying repository.

1. PHASE MODULATION AND THE COUPLING PARAMETER

Chowning's foundational synthesis technique [1] is conventionally termed frequency modulation (FM) but is implemented as phase modulation (PM): the carrier phase is offset by a scaled modulator, $y(t) = \sin\big(\omega_c t + I\sin(\omega_m t)\big)$, with carrier and modulator angular frequencies ω_c and ω_m. For one-way sinusoidal FM the modulation index is $I = \Delta\omega/\omega_m$, the ratio of peak frequency deviation to modulator frequency; in the equivalent PM form, the same quantity appears directly as the peak phase deviation in radians [2]. With the index so matched, the two formulations are spectrally equivalent in this topology; we adopt the PM form throughout.

Copyright: © 2026. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 Unported License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

In cross-modulation no carrier/modulator hierarchy exists: both oscillators modulate each other simultaneously. Because I is tied to a designated modulator frequency, which has no analogue in the symmetric case, we define the coupling directly by the peak phase deviation $\beta \in \mathbb{R}$: a frequency-independent gain, equal in both coupling directions (asymmetric coupling is a natural extension). The discrete-time PM oscillator advancing at angular frequency ω with sample period T_s is

$$x_n = \sin(\phi_n + \beta \cdot m_n), \qquad \phi_n = \phi_{n-1} + \omega T_s, \quad (1)$$

where m_n is the modulation signal.

2. FM CROSS-MODULATION

Cross-modulation, the simultaneous mutual frequency or phase modulation of two oscillators, has a long history in analog and digital synthesis as a source of complex, chaotic timbres. Slater's account of cross-coupled Buchla oscillators [3] documents the system's capacity to produce periodic, quasiperiodic and chaotic sounds as coupling is increased, describing it as a simple and powerful generator of a wide range of complex timbres. Today, numerous commercially available dual-core VCOs are designed to allow for cross-coupled FM. Valsamakis and Miranda [4] treat the digital version of cross-coupled FM synthesis explicitly as an iterated nonlinear map, exploring its chaotic regimes for compositional purposes. These two are not quite the same phenomenon: as we show below, the "chaos" of the digital map depends on how the per-sample loop is resolved — and, for two oscillators, is not inherited from the continuous system at all.

In PM terms, cross-modulation couples two oscillators so that each modulates the phase of the other with the current output sample. Substituting (1) into itself for both oscillators yields the *implicit* system

$$x = \sin(\phi_c + \beta \cdot y), \qquad y = \sin(\phi_m + \beta \cdot x), \quad (2)$$

where ϕ_c and ϕ_m advance independently each sample. Since x depends on y and y depends on x, neither can be computed directly: the system forms an algebraic loop.

By the Brouwer fixed-point theorem, a solution to (2) always exists for any β, since the map $G : [-1,1]^2 \to [-1,1]^2$ defined by the right-hand sides maps a compact convex set into itself. Existence, however, does not imply uniqueness: G is a contraction only for $\beta < 1$ (Section 5), and above this boundary multiple solutions may coexist for

a given phase pair. "The" solution of (2) is then well defined only relative to a selection rule. The iterative solvers compared below, which start from the previous sample's solution, select a branch as a by-product of where their iterations settle — typically the branch continuous with the previous sample; the fourth realisation states the rule outright, selecting the closest solution at every sample. The choice of method thus determines which solution is selected at each sample, making it an important musical and sound design choice.

This paper contributes (i) a formulation of digital cross-modulation as a per-sample implicit system, with an explicit account of solution existence, uniqueness, and branch selection; (ii) a unified view of the delayed, iterative, averaged, and root-tracking realisations as one family of per-sample iterations, each started from the previous sample's solution; (iii) a comparative dynamical and spectral analysis of the four realisations at native rate and under $64\times$ oversampling; and (iv) open-source real-time implementations.

3. STANDARD IMPLEMENTATION: ONE-SAMPLE DELAY

The standard digital resolution of the algebraic loop substitutes the previous sample's output as the modulation argument, replacing the implicit system (2) with the explicit recurrence

$$x_n = \sin(\phi_{c,n} + \beta \cdot y_{n-1}), \qquad y_n = \sin(\phi_{m,n} + \beta \cdot x_{n-1}). \tag{3}$$

This is a *two-dimensional explicit recurrence map* on the state space $[-1, 1]^2$: the state (x_n, y_n) is determined entirely by the previous state (x_{n-1}, y_{n-1}) and the current phase pair $(\phi_{c,n}, \phi_{m,n})$, which act as a quasi-periodic external drive. The map is iterated once per sample; its long-run dynamics — fixed points, periodic orbits, strange attractors — constitute the synthesis method's sonic behaviour.

The relation between (3) and the implicit system (2) is regime-dependent. Where the coupling is weak and the signals vary slowly relative to the sample rate, $y_{n-1} \approx y_n$, and the delayed recurrence is a consistent approximation of (2) whose error shrinks with the ratio of signal bandwidth to sample rate. In the strongly coupled regime, however, the delay acts as structural memory that changes the fixed-point structure, the stability boundaries, and the attractor geometry: at practical sample rates, (3) then behaves as a dynamical system in its own right rather than as an approximation of (2). As we show below, the difference is qualitative across much of the musically relevant range of β.

The delay can be placed in three ways: on both coupling paths, as in (3), or on a single path, delaying only x or only y. We analyse the symmetric two-delay form because it preserves the exchange symmetry $x \leftrightarrow y$ of (2) and corresponds to the parallel evaluation of two oscillator nodes from the previous tick's outputs, as in graph-based audio environments. The single-delay forms are sequential updates that break this symmetry — and pose the question of which oscillator to privilege; as Section 5 makes explicit, each of them coincides with a single sweep — one pass through both updates — of the fixed-point iteration (4) in the corresponding update order.

4. ZERO-DELAY FEEDBACK

Zavalishin's topology-preserving transform (TPT) framework — commonly referred to as zero-delay feedback (ZDF) — addresses exactly this problem in the context of virtual analog filter design [5]. The central observation is that analog filter topologies contain no internal delays: the circuit responds to its current state instantaneously. Inserting a unit delay to resolve the digital feedback loop introduces an artificial discrepancy between the analog prototype and its digital model, manifesting as tuning errors, instability, and incorrect resonance behaviour.

Mathematically, resolving a delay-free loop amounts to solving a system of implicit equations at every sample — a problem treated in the discrete-time modelling literature well before the ZDF terminology emerged, notably in the K-method of Borin, De Poli and Rocchesso [6] and its extensions to nonlinear filter networks [7]. Zavalishin's framework recasts this practice for virtual analog design: retain the implicit equation and solve it algebraically (for linear systems) or numerically (for nonlinear systems) at each sample, preserving the topology of the analog prototype. The approach has been widely adopted for ladder and state-variable filters, but to our knowledge neither ZDF nor the older delay-free-loop methods have previously been applied to oscillator cross-modulation.

The cross-modulation system (2) has the same algebraic loop structure as a nonlinear filter: a bounded nonlinearity ($\sin$) in a feedback path. The same per-sample implicit formulation applies directly. Unlike the filter case, the bounded range of $\sin$ guarantees that a solution always exists and that iterative solvers cannot diverge to infinity.

5. COMPARED METHODS

We compare four realisations of cross-modulation. The first three — the one-sample delay, FPI, and the trapezoidal Newton scheme — are evaluated at both the native sample rate ($f_s = 48\,\text{kHz}$) and with $64\times$ oversampling ($\approx 3.07\,\text{MHz}$, decimated via a linear-phase FIR anti-aliasing filter); the fourth, a root-tracking solver described at the end of this section, is evaluated at the native rate only. Frequencies are set to $f_c = 200\,\text{Hz}$ and $f_m = f_c \cdot \varphi$ where $\varphi = (1+\sqrt{5})/2$ is the golden ratio. The two oscillator phases (ϕ_c, ϕ_m) constitute a two-dimensional periodic state space that is topologically a torus. With the irrational frequency ratio — as with φ — the joint phase sequence never repeats, so the drive never settles into a periodic cycle.

1-sample delay (Delayed). System (3) is the universal default in block-based DSP environments and serves as the baseline. Its dynamics are governed by an explicit 2D map that differs structurally from the implicit system.

Fixed-point iteration (FPI). The implicit system (2) is solved per sample by alternating iteration

$$y^{(k+1)} = \sin(\phi_m + \beta\, x^{(k)}), \quad x^{(k+1)} = \sin(\phi_c + \beta\, y^{(k+1)}), \tag{4}$$

initialised from the previous sample (x_{n-1}, y_{n-1}) and run for $N = 16$ iterations. A single sweep ($N = 1$) reproduces exactly the single-delay realisation of Section 3 with the delay on the $x \to y$ path; the opposite update order yields the other single-delay variant. Seen this way, the explicit and iterative realisations form a single family of per-sample iterations, each starting from the previous sample's solution: the two-delay map is a one-step simultaneous update, the single-delay maps are one-step sequential updates, FPI raises the sweep count towards the solution of (2) wherever the iteration contracts, and the root-tracking solver (below) finds the solutions of (2) directly, whether or not the iteration contracts. By the Banach fixed-point theorem, iteration (4) converges to the unique solution of (2) for $\beta < 1$. Above $\beta = 1$ contraction is no longer guaranteed; empirically the iteration then fails to settle and moves among coexisting solution branches, so that FPI with fixed N defines a synthesis map of its own rather than a solver for (2).

Trapezoidal Newton (Trap). Both the delayed system (3) and FPI (4) evaluate the modulation argument at a single time point — either the previous sample or the current one. Because the coupling in (2) is algebraic rather than an integrator structure, the trapezoidal rule of ZDF filter design does not transfer here as the discretisation of an underlying differential equation. We adopt it instead as a design principle: following the trapezoidal update of topology-preserving filter design, we replace the instantaneous coupling argument $\beta \cdot y_n$ with the two-point average $\frac{\beta}{2}(y_n + y_{n-1})$ and advance the phase to the half-step midpoint. The resulting implicit averaged system is

$$\begin{aligned} x_n &= \sin\Big(\phi_{c,n-\frac{1}{2}} + \tfrac{\beta}{2}(y_n + y_{n-1})\Big), \\ y_n &= \sin\Big(\phi_{m,n-\frac{1}{2}} + \tfrac{\beta}{2}(x_n + x_{n-1})\Big), \end{aligned} \tag{5}$$

where $\phi_{c,n-1/2} = \phi_{c,n-1} + \Delta\phi_c/2$ is the half-step phase. The averaging has two consequences. First, it acts as a one-zero lowpass on the modulation path (with a zero at Nyquist): system (5) mildly filters the modulator and is therefore itself a modification of (2), not a transparent solver for it. Second, the coupling seen by each unknown is $\beta/2$ rather than β, which extends the guaranteed contraction boundary from $\beta < 1$ to $\beta < 2$; averaging over more past samples would extend the boundary further, at the cost of filtering the modulation path more heavily.

Because system (5) is nonlinear, we solve it per sample with Newton-Raphson iteration, initialised — like FPI — from the previous sample's solution (x_{n-1}, y_{n-1}). At each iteration, the current guess (x, y) is substituted into both sides of (5); the mismatch is used to compute a correction step, which can be derived in closed form from the derivatives of $\sin$. We limit each correction step by $\Delta x \leftarrow r \tanh(\Delta x / r)$ with $r = 0.5$ to prevent overshooting. For $\beta < 2$ the correction step is well defined at every iteration — the quantity by which it divides is bounded away from zero — so the safeguard that reverts to the previous sample's values on undefined output is relevant only when the system is driven beyond the contraction boundary. Across all experiments reported here, the safeguard never activated for $\beta < 2$ (2.6×10^8 samples rendered); in 30 s stress renders at $\beta \in \{2.5, 3.0\}$ it likewise never activated — the step limiting keeps every iterate finite even beyond the boundary. For $\beta < 2$ the averaged system thus has a unique solution; the safeguarded Newton solver converged to it to machine precision within $N = 8$ iterations in every tested case.

Root tracking (Tracked). The fourth realisation solves (2) itself at every sample: its output satisfies the stated equations, to machine precision, for every β — not the equations of a modified system. Substituting the second equation into the first reduces the system to a single equation $g(x) = x - \sin\big(\phi_c + \beta \sin(\phi_m + \beta x)\big) = 0$, whose solutions necessarily lie in $[-1, 1]$, where $g(-1) \le 0 \le g(1)$. A dense scan of this interval followed by bisection on every sign change locates the coexisting solutions to machine precision, without derivatives and without the possibility of undefined output; root pairs closer than the scan spacing, or points where g touches zero without crossing, can momentarily escape detection. Among the solutions found we select the one closest to the previous sample's, tracking a branch by continuity. Unlike the iterative methods, which follow only the branch they are pulled toward, this full scan can hold a branch the others lose. Where the iteration contracts it is the limit FPI approaches; beyond, it realises one branch among several. The same scan yields the number of detected solutions at each sample. Continuity is itself a selection rule, not a neutral fact: the same scan could equally select the smallest solution, or hold branches with hysteresis, and above $\beta = 1$ each rule yields a different signal, every one of which satisfies (2). System (2) thus defines not one canonical signal but a family of them; root tracking realises the member given by minimal discontinuity, jumping only where its branch disappears. For $\beta < 1$, where the solution is unique, all rules coincide. In the real-time implementation the full scan is not repeated at every sample. Newton's method refines the previous root, and a short sweep of the surrounding interval — using the bound $|g'| \le 1 + \beta^2$, which limits how close an undetected solution can lie — confirms that no other solution is closer to the previous sample's; only when this check fails, or the tracked branch disappears, is the full scan rerun. Across all experiments reported here this produces output sample-identical to the per-sample scan, at a cost about that of the iterative solvers at $\beta = 1.4$ ($\approx 1.5\times$ in the worst case at $\beta = 3.5$); it runs comfortably in real time and can itself be oversampled.

6. DYNAMICAL ANALYSIS

All results are presented in two conditions: at the native sample rate of 48 kHz and with $64\times$ oversampling. Com-

paring these conditions shows that, at the factor tested, oversampling removes the delayed system's most audible artefact and substantially reduces — but does not eliminate — its deviation from the continuity-tracked realisation of (2) in the strongly coupled regime.

6.1 Bifurcation Structure

Fig. 1 shows bifurcation diagrams at 48 kHz. All methods agree in the low-β quasiperiodic regime. Above $\beta \approx 1$ the methods diverge qualitatively. The delayed system enters a broad, structureless regime at $\beta \approx 1.4$ (though, as Section 7 shows, positive Lyapunov exponents appear only at higher β). Plain FPI shows organised curved bands from $\beta \approx 1.0$ to $\beta \approx 1.6$; above this range FPI cycles among multiple fixed points rather than converging, defining a distinct dynamical system. Trapezoidal Newton exhibits the most organised structure, with clean branches up to $\beta \approx 2.5$. The root-tracking panel shows organised curved bands over the whole range, emerging continuously from the low-β quasiperiodic regime with no chaotic scatter; among the realisations, trapezoidal Newton follows its band geometry most closely, with visually coincident structure up to $\beta \approx 2.5$, beyond which the trapezoidal panel adds noisy vertical bursts absent from the tracked solution. A dense scan of (2) along the trajectory confirms the multiplicity picture: a single solution exists at every sample for $\beta < 1$, while the scan finds up to 3 solutions at $\beta = 1.4$ and up to 7 at $\beta = 3.5$.

The corresponding bifurcation diagrams at $64\times$ oversampling (available in the repository) show that oversampling smooths the delayed system's transition to chaos but does not recover the branch structure of the implicit solvers. The trapezoidal solver is largely unchanged under oversampling, consistent with its averaged coupling and half-step phase differing from (2) only at second order in the sample period.

6.2 Spectral Transitions

Fig. 2 shows spectrograms at 48 kHz (audio examples 1–4). The delayed system's $f_s/4$ artefact is visible as a persistent bright band at ≈ 12 kHz from $\beta \approx 1.2$ onward (audio example 5). Its location follows from the structure of the delayed coupling: with a delay on both paths, a perturbation of x returns to x only after two samples, so the map contains an effective two-sample feedback loop. When the round trip is strong and sign-alternating, this loop supports near-period-4 behaviour — energy concentrated near $f_s/4$ — much as a one-sample negative feedback loop rings at $f_s/2$. The single-delay variant behaves analogously, with its artefact band near $f_s/2$ rather than $f_s/4$, consistent with its one-sample round trip. The diagram also exhibits abrupt transitions, consistent with bifurcations of the delayed map, visible as discontinuous shifts in the pre-chaotic periodic structure; these particular events are features of the delay map and are absent from the plotted implicit realisations. Plain FPI shows a continuous broadband transition without abrupt spectral events. Trapezoidal Newton shows the latest and most gradual onset of spectral broadening, consistent with its extended convergent regime. The root-tracking solver is close to trapezoidal Newton but free of the delay map's $f_s/4$ band and of its abrupt transitions, giving the smoothest spectral evolution of the four.

At $64\times$ oversampling, the delayed system's $f_s/4$ artefact vanishes (audio example 6) — the internal sample rate shifts $f_s/4$ to ≈ 768 kHz, far above the decimated Nyquist frequency. The oversampled spectrograms of FPI and trapezoidal Newton are visually similar to their 48 kHz counterparts: at the tested oversampling factor, the artefact is corrected but the three oversampled realisations remain dynamically and spectrally distinct.

6.3 Time-Domain Waveforms

Fig. 3 shows 10 ms time-domain excerpts at 48 kHz. At $\beta = 1.4$ the delayed system produces high-frequency oscillations between successive samples, while the trapezoidal solver and the root-tracking solver trace a smooth waveform, nearly indistinguishable from one another. Plain FPI remains visually smooth at $\beta = 1.4$, although beyond $\beta = 1$ it no longer settles on a solution of (2) (Section 5); at $\beta = 1.8$ its cycling behaviour produces a waveform distinct from both the delayed and the trapezoidal systems (audio example 7).

6.4 Phase Portraits

Fig. 4 shows phase portraits at 48 kHz. At $\beta = 0.8$, all four methods trace similar quasi-periodic curves. At $\beta = 1.4$, the attractors diverge: the delayed system's portrait is more diffuse, plain FPI traces a reduced set of curves reflecting its cycling behaviour, and the trapezoidal solver and the root-tracking solver both maintain well-defined orbit structure. At $\beta = 2.0$, the delayed system fills the state space nearly uniformly — dense, though as Table 1 shows just short of the onset of chaos — while plain FPI shows dense scatter from its non-convergent iterations. The trapezoidal solver and the root-tracking solver still exhibit coherent attractor geometry, nearly coincident, consistent with the averaged system (5) remaining close to (2) at this rate. The corresponding $64\times$ oversampled phase portraits and time-domain waveforms are available in the repository.

7. SENSITIVE DEPENDENCE AND THE ORIGIN OF CHAOS

Table 1 reports the largest conditional Lyapunov exponent of each realisation: both trajectories receive the identical phase sequence, and the exponent measures whether a small perturbation of the output state grows or decays under that common drive — positive when nearby trajectories separate exponentially. For root tracking, no positive exponent was observed at any tested β: the phases advance at fixed rates (a rigid quasiperiodic rotation), and the branch memory can change which solution is selected — a discontinuous, finite-size effect — but produced no exponential separation in our measurements. Its signal is quasiperiodic, with jumps only where a branch disappears; the trapezoidal solver likewise stays contracting ($\lambda < 0$) across the range tested. Positive exponents appear only in the delayed and

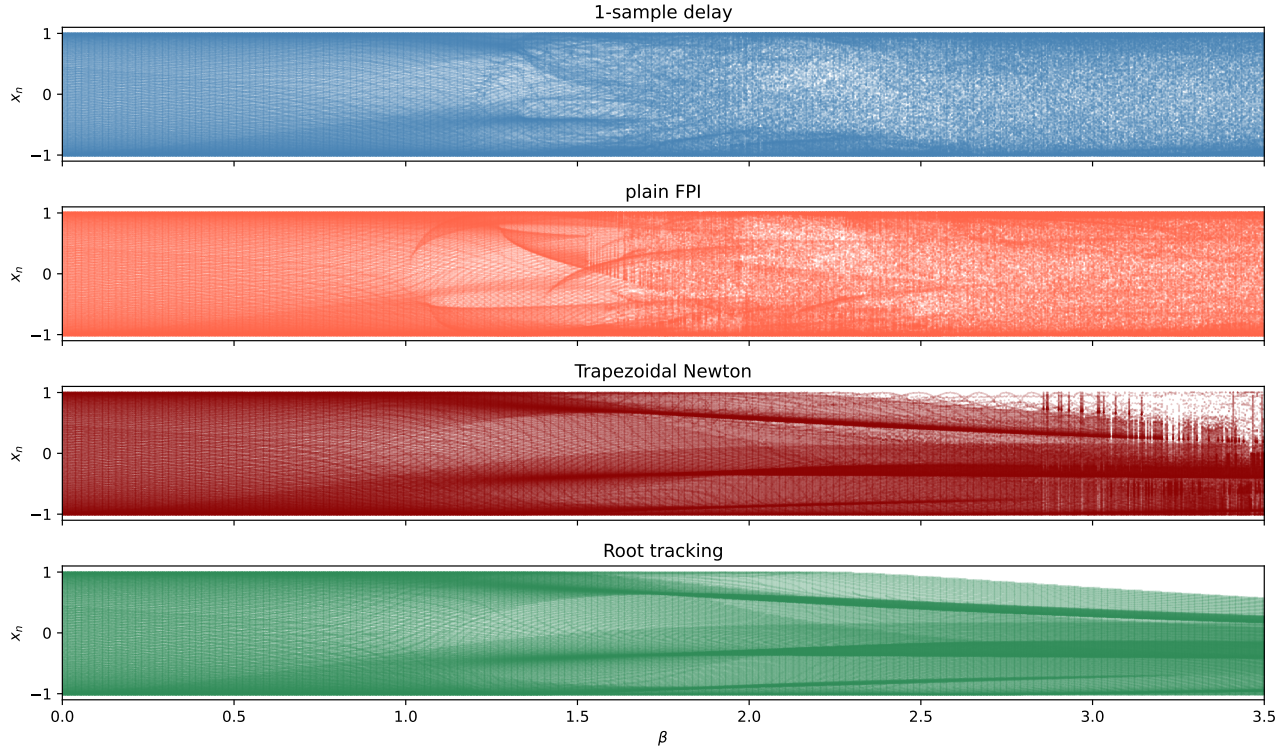


Figure 1. Bifurcation diagrams at 48 kHz for the four realisations, including the tracked solution of (2), over $\beta \in [0, 3.5]$, $\omega_m/\omega_c = \varphi$. Each panel plots the long-run values of x_n after discarding transients.

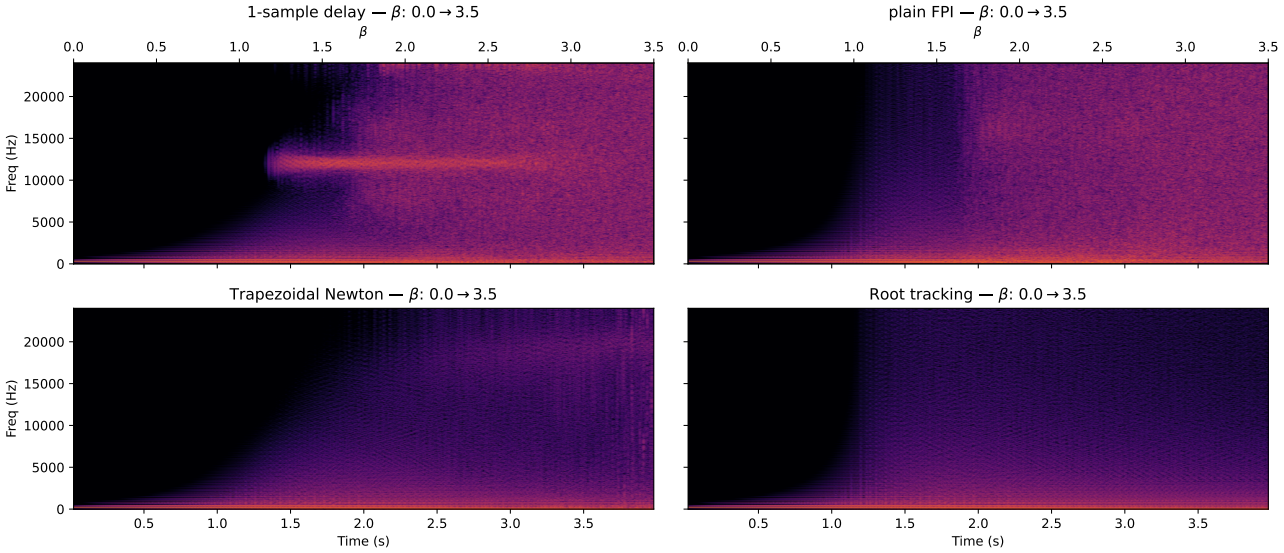


Figure 2. Spectrograms of a β-sweep from 0 to 3.5 over four seconds at 48 kHz, $\omega_m/\omega_c = \varphi$, for the four realisations (top: delayed, FPI; bottom: trapezoidal Newton, root tracking).

fixed-N FPI maps, and only at high β (from $\beta \approx 2.3$ and 2.0 respectively): there the one-sample delay, or the truncated iteration, feeds each output back into the next state, and it is this memory — absent from (2) itself — that produces the sensitive dependence. The broadband structure at lower β is thus not chaos but a complex non-chaotic attractor shaped by the quasiperiodic drive.

This sharpens the relation to analog cross-modulation, long heard as chaotic [3]. The continuous counterpart of two cross-coupled oscillators — coupling through frequency, $\dot{\phi}_1 = \omega_1 + \beta \sin \phi_2$, $\dot{\phi}_2 = \omega_2 + \beta \sin \phi_1$ — is an autonomous flow on the two-dimensional phase torus, and such flows cannot be chaotic: sensitive dependence in continuous time requires at least three dimensions [8]. Numerical integration of the pair confirms quasiperiodic or locked behaviour throughout. Where analog cross-modulation is

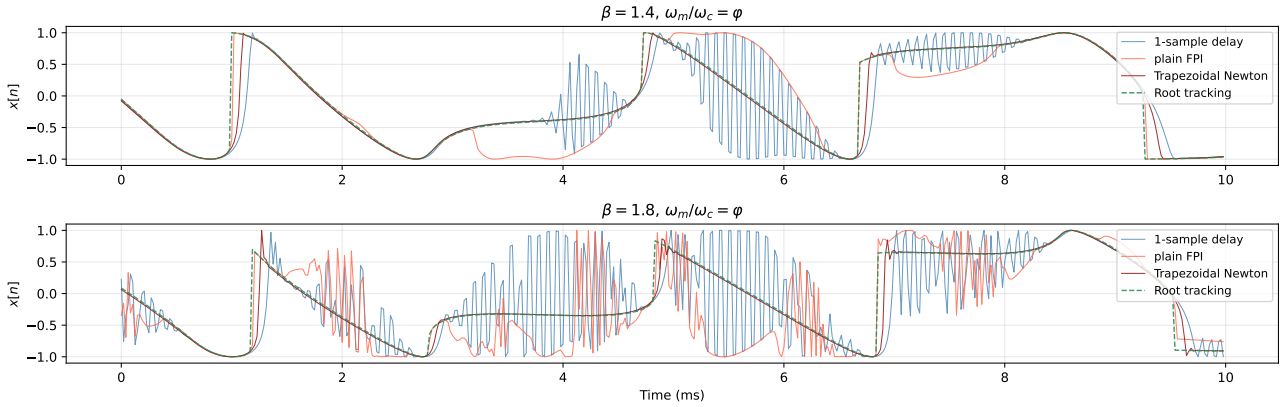


Figure 3. Time-domain waveforms (10 ms excerpt) at 48 kHz for $\beta = 1.4$ and $\beta = 1.8$, $\omega_m/\omega_c = \varphi$. The delayed system exhibits large inter-sample discontinuities absent from the implicit solvers; the tracked solution (dashed) nearly coincides with trapezoidal Newton.

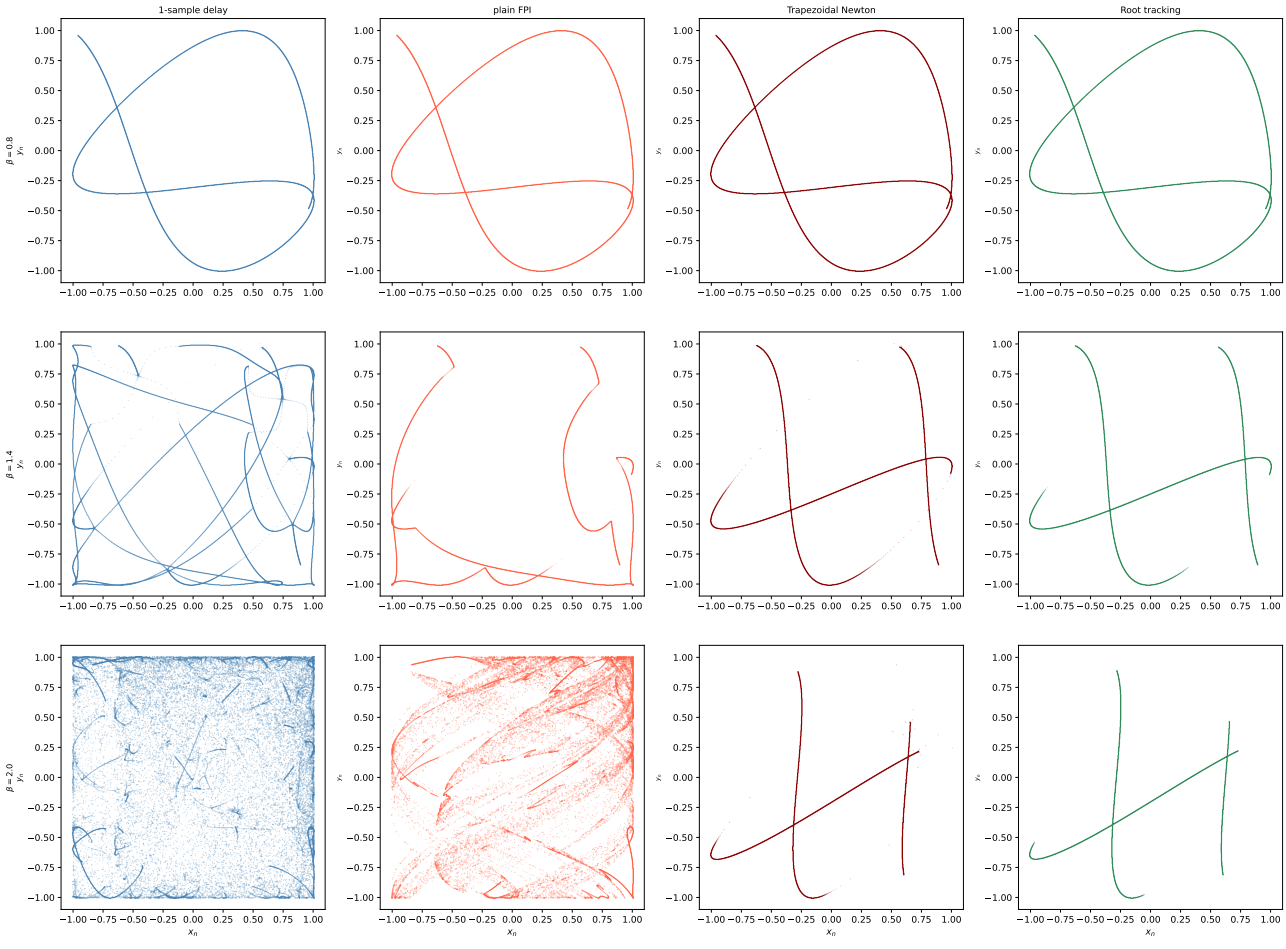


Figure 4. Phase portraits (x_n vs. y_n) at 48 kHz for $\beta \in \{0.8, 1.4, 2.0\}$ (columns: delayed, FPI, trapezoidal Newton, root tracking). At $\beta = 0.8$ all four trace similar torus windings. At $\beta = 2.0$ the delayed system fills the state space uniformly while the trapezoidal solver and the root-tracking solver maintain organised orbit structure.

genuinely chaotic, the extra dimensions come from the hardware — oscillator cores, converters, filters — or from more than two oscillators; in preliminary experiments, three digital oscillators in a ring suffice. For two oscillators, then, sensitive dependence arises only through the memory a discrete resolution introduces: the delayed and fixed-N FPI maps turn chaotic where the continuity-tracked solution of (2) and its continuous counterpart show none (the latter necessarily so).

Table 1. Largest conditional Lyapunov exponent λ per sample of the output state, estimated by two-trajectory renormalisation with identical phase sequences in both trajectories; $\lambda > 0$ indicates sensitive dependence (chaos). For root tracking, 0 denotes that no exponential separation was observed. Positive exponents occur only for the delayed and fixed-N FPI maps, and only at high β.

β	delayed	FPI	trap.	root tr.
0.8	-0.9	-11.1	-1.5	0
1.4	-0.4	-7.5	-0.9	0
2.0	-0.1	$+0.3$	-0.6	0
2.5	$+0.1$	$+3.4$	-0.4	0
3.5	$+0.5$	$+13.7$	-0.1	0

8. IMPLEMENTATION

All four realisations are implemented in C++ as a single DSP core shared by a SuperCollider unit generator (UGen) and a JUCE-based VST3/AU audio plugin, cross-validated against the Python implementations. The realisation and iteration count are switchable at run time. On a Ryzen AI 9 365, the delayed, FPI ($N = 16$) and trapezoidal Newton ($N = 8$) realisations render at approximately 31 ns, 440 ns, and 430 ns per sample (670×, 48× and 49× real time at 48 kHz, single core); the root-tracking solver costs roughly 440 ns per sample at $\beta = 1.4$ ($\approx$ 660 ns in the worst case at $\beta = 3.5$), around 48× real time — comparable to the iterative solvers, its cost varying with β as branches appear and disappear more often. The trapezoidal Newton solver uses the step limiting and fallback described in Section 5. A minimal SuperCollider usage example is included in the repository. All figures in this paper were generated by a Python script using Numba-accelerated implementations of each realisation. Source code, the JUCE plugin, the SuperCollider UGen, the Python analysis scripts, and the numbered audio examples referenced in Section 6 are available at `https://github.com/lucdoebereiner/ZDFFMSMC`.

9. CONCLUSION

In the weakly coupled regime, the one-sample-delay implementation of FM cross-modulation approximates the implicit coupled system. Beyond it, the delay is no longer an approximation: it defines a structurally different dynamical system, producing artefacts and dynamical regimes with no counterpart in the implicit system (2). Oversampling removes the most audible of these artefacts and, at the factors tested, substantially reduces the delayed system's deviation from the continuity-tracked realisation of (2): at $\beta = 1.4$, the band-averaged spectral deviation (mean absolute difference of the Welch log-power spectra, 20 Hz–20 kHz) fell from 9.0 dB at native rate to 1.9 dB at 64×. It remained nonzero at the largest factor tested; whether it vanishes asymptotically is left open, and the comparison is specific to the continuity rule. The full oversampling study is in the repository. Trapezoidal Newton-Raphson solves its averaged implicit system (5) — contractive, with a unique solution, for $\beta < 2$ — to machine precision in eight iterations throughout the tested range; and the root-tracking solver, at modest extra cost, makes solutions of (2) itself — under an explicit branch-selection rule — playable in real time. The choice of realisation is not an engineering detail but a compositional one: each method defines a distinct dynamical system with its own bifurcation structure and spectral character. Future work includes asymmetric coupling $\beta_1 \neq \beta_2$; the true-FM formulation, in which the coupling modulates frequency, the pair becomes two differential equations — quasiperiodic for two oscillators (Section 7) — and trapezoidal integration applies in its native sense; and networks of three or more oscillators, where continuous-time chaos first becomes possible.

Acknowledgments

The authors thank the anonymous reviewers for their detailed and constructive comments.

10. REFERENCES

[1] J. M. Chowning, "The synthesis of complex audio spectra by means of frequency modulation," *Journal of the Audio Engineering Society*, vol. 21, no. 7, pp. 526–534, 1973.

[2] K. Nielsen, "Practical linear and exponential frequency modulation for digital music synthesis," in *Proceedings of the 23rd International Conference on Digital Audio Effects (DAFx-20)*, Vienna, Austria, 2020, pp. 132–139.

[3] D. Slater, "Chaotic sound synthesis," *Computer Music Journal*, vol. 22, no. 2, pp. 12–19, 1998.

[4] N. Valsamakis and E. R. Miranda, "Iterative sound synthesis by means of cross-coupled digital oscillators," *Digital Creativity*, vol. 16, no. 2, pp. 90–98, 2005.

[5] V. Zavalishin, *The Art of VA Filter Design*. Native Instruments, 2012, revised edition 2018. Available at https://www.native-instruments.com/fileadmin/ni_media/downloads/pdf/VAFilterDesign_2.1.0.pdf.

[6] G. Borin, G. De Poli, and D. Rocchesso, "Elimination of delay-free loops in discrete-time models of nonlinear acoustic systems," *IEEE Transactions on Speech and Audio Processing*, vol. 8, no. 5, pp. 597–605, 2000.

[7] F. Fontana and F. Avanzini, "Computation of delay-free nonlinear digital filter networks: Application to chaotic circuits and intracellular signal transduction," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4703–4715, 2008.

[8] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, 2nd ed. Boulder, CO: Westview Press, 2015.