

Advanced Encoding and Rewritability Techniques for Coherent Wave Memory

Mike Tierce

Independent Researcher

ORCID: 0009-0004-3869-958X | github.com/miketierce/cwm | Version

Abstract—

Coherent Wave Memory (CWM) encodes information in the acoustic eigenmode spectrum of glass resonators and computes via wave interference [6]. This companion paper develops twenty-two modeled extensions to the baseline CWM architecture. Six firmware-level techniques—synaptic pruning (+10.7% recall accuracy), in-situ Boolean computation (>90% fidelity), mode hybridization (+160% capacity), null-space multiplexing (+60%), polysemic readout (+297%), and phase-spectral encoding (+84% discriminability)—require zero hardware changes. Sixteen cross-

domain investigations test 87 additional hypotheses across wave physics, information theory, and spectral analysis, confirming 51 and killing 36. Three paths to rewritability—firmware-defined virtual rewriting (4+ logical devices per rod), binary MEMS perturbation sites (7.6 bits, <0.5% Q penalty), and writable shell coatings (100 nm at $Q > 5,000$)—progressively transform CWM from read-only to fully reconfigurable. All extensions are simulated and validated by automated tests; none have been confirmed on hardware.

PART I

Theory and Architecture

1. INTRODUCTION

Coherent Wave Memory (CWM) stores data in the eigenmode spectrum of solid glass resonators and performs nearest-neighbor search via acoustic wave interference in a single propagation cycle. The core architecture—eigenmode encoding, Rayleigh perturbation writing, and interference recall—is developed and validated at macro scale in the companion paper [6], which establishes a 230prototype(38 core materials) with 98.5 dB derived SNR, 9,380 thermally stable modes per resonator, and a five-mechanism MEMS Q-factor model predicting $Q_{\text{total}} = 9,097$. Full details on the architecture (eigenmode encoding, Shannon capacity, Hopfield associative recall), substrate selection, finite element validation, MEMS scaling laws, Q-factor analysis, device specification, and fabrication pathway are in [6].

This paper extends the baseline architecture in two directions. First, Section 2 presents six firmware-implementable encoding and recall tech-

niques that enhance capacity and functionality with zero hardware changes, followed by sixteen cross-domain investigations that systematically test the boundaries of eigenmode physics across wave mechanics, information theory, and spectral analysis. Second, Section 3 develops three hardware paths to rewritability—firmware-defined virtual rewriting, binary MEMS perturbation sites, and writable shell coatings—progressively transforming CWM from a glass harmonica (fixed pitch) to a glass armonica (reconfigurable).

All extensions are modeled by simulation (48 modules total; 1,747 tests across 20 modules for cross-domain investigations alone) and await hardware validation on the MEMS prototype described in [6].

2. ADVANCED ENCODING AND RECALL TECHNIQUES

Note. The techniques in this section are modeled extensions built on the core CWM formalism [6]. All results are from simulation; none have been validated on physical hardware.

The core CWM architecture [6] establishes the fundamental device physics. This section presents twenty-two modeled extensions—discovered through systematic computational exploration of the eigenmode physics—that could extend CWM’s capabilities beyond the baseline architecture if confirmed on hardware. The first six (§§2.1–2.6) require **zero hardware changes**: they would be implementable as firmware-level signal processing on the CMOS readout die. The remaining sixteen explore broader extensions through cross-domain physical and mathematical analysis. Section 2.7 summarises these investigations in tabular form; full details are available in the companion simulation modules and companion volume [5].

2.1 Synaptic Pruning for Associative Recall

The problem. CWM’s associative recall [6, §2.3] is mathematically equivalent to a Hopfield network, where the “weight matrix” is the set of stored spectral fingerprints. As the number of stored patterns P approaches the capacity limit ($P_{\text{max}} \approx 0.138 N$ for < 1% bit-error rate [1], where N is the number of modes), something goes wrong. Each pattern’s fingerprint is a vector of N mode amplitudes, and storing many patterns in the same weight matrix creates inter-pattern crosstalk: the fingerprint of pattern A partially overlaps with patterns B, C, D, etc. When you query for pattern A, the response includes “ghost” contributions from these other patterns, which can push the recall toward a spurious attractor—a false match.

Think of it like overhearing multiple conversations in a crowded room. Each conversation (pattern) is carried by the same physical medium (the air / the mode spectrum). When only a few people are talking, you can follow any one conversation clearly. When the room is full, the conversations

blur together and you mishear words. The question is: can you improve your hearing without changing the room?

The approach. The inter-pattern crosstalk is concentrated in the *small-magnitude* entries of the weight matrix—the weak, non-specific couplings that correlate with multiple patterns rather than encoding any single one. We hypothesised that zeroing these small weights—a form of controlled “forgetting”—could remove crosstalk noise while preserving the strong weights that encode the actual patterns.

This is directly analogous to synaptic pruning in biological neural development [2]. During brain maturation, the nervous system eliminates roughly 50% of its synapses between early childhood and adulthood. Far from being a deficiency, this pruning improves signal-to-noise ratio by removing weak, non-specific connections that add noise to neural circuits. The mature brain is more capable than the infant brain, with fewer synapses.

The experiment. We store $P = 8$ binary patterns in an $N = 50$ Hopfield network (load factor $P/N = 0.16$, well within the overload regime where crosstalk matters) and measure recall accuracy under noisy queries (20% of bits randomly flipped). We then sweep a pruning threshold θ : all weight-matrix entries with magnitude $|w_{ij}| < \theta$ are zeroed. For each threshold, we run 40 independent trials with random noise realizations and average the recall accuracy.

Results. Without pruning ($\theta = 0$): recall accuracy 0.700 (70.0%)—the network gets the right answer 70% of the time. With optimal pruning at $\theta^* = 0.055$ (corresponding to zeroing all weights below 5.5% of the maximum weight magnitude): **0.775 (77.5%)**, a gain of **+10.7%**.

Pruning threshold θ	Recall accuracy	Change vs. baseline
0 (no pruning)	0.700	—
0.027	0.725	+3.6%
0.055	0.775	+10.7%
0.082	0.750	+7.1%
0.110	0.700	0%
0.190	0.575	-17.9%
0.300	0.475	-32.1%

The optimum is sharp. Below $\theta^* = 0.055$, pruning removes too few weights to matter. Above it, pruning begins cutting into the pattern-encoding weights themselves, destroying signal along with noise. The optimal threshold removes approximately 40% of all weight-matrix entries—the smallest 40%—which turns out to be almost entirely inter-pattern crosstalk.

What this means for CWM. In CWM’s associative recall mode, the read-out ASIC computes a correlation score between the query spectrum and each rod’s stored spectrum. Pruning corresponds to applying a spectral mask: before computing the correlation, zero out all frequency components whose amplitude falls below a threshold. This is a single line of firmware—a thresholded multiply—applied to the FFT output. At high pattern loads (hundreds of patterns per rod, approaching the Hopfield capacity limit), this mask recovers 10.7% of the recall accuracy lost to inter-pattern crosstalk.




Figure 1. (a) Synaptic pruning concept: zeroing small-magnitude weights removes inter-pattern crosstalk. (b) Recall accuracy vs. pruning threshold, showing +10.7% gain at optimal $\theta = 0.055$.

2.2 In-Situ Boolean Computation

The insight. CWM’s modes are linear oscillators, and wave superposition is linear. If two patterns A and B are encoded as mode amplitudes in the same rod (or superposed from two rods’ responses), the combined amplitude at each mode frequency is the *sum* of the individual amplitudes. This summed signal contains enough information to extract the Boolean functions AND, OR, and XOR of the two binary patterns—without a separate compute step.

Why this works. Consider a single mode where pattern A encodes a ‘1’ (high amplitude, $a = 1.0$) and pattern B encodes a ‘0’ (low amplitude, $b = 0.2$). The combined amplitude is $a + b = 1.2$. Now consider the four possible bit combinations:

A bit	B bit	Combined amplitude	AND	OR	XOR
0	0	0.4	0	0	0
0	1	1.2	0	1	1
1	0	1.2	0	1	1
1	1	2.0	1	1	0

The combined amplitudes cluster into three groups: low (0.4), medium (1.2), and high (2.0). Each Boolean function corresponds to a different partitioning of these groups:

- **AND:** only the high group (both bits = 1) → threshold at 70% of max
- **OR:** medium and high groups (at least one bit = 1) → threshold at 30% of max
- **XOR:** only the medium group (exactly one bit = 1) → band-pass between 30% and 75% of max

The experiment. We encode two random binary patterns A and B as mode amplitudes in a 32-mode system (high = 1.0, low = 0.2—the low value is non-zero to maintain mode excitation), evolve each pattern’s wave representation through $Q = 800$ oscillation cycles, superpose the resulting signals, and decode the combined amplitudes using the three threshold rules above.

Results.

Operation	Fidelity	Method
XOR	90.6%	Band-pass: 30%–75% of max amplitude
AND	96.9%	High-pass: >70% of max amplitude
OR	93.8%	Low-pass: >30% of max amplitude

All three operations exceed 90% fidelity from a **single readout cycle** under ideal simulation conditions (no phase noise, perfect frequency knowledge). Real-device fidelity will depend on readout SNR and mode-tracking accuracy; the threshold decoding is intrinsically robust to additive noise but sensitive to systematic frequency drift. The conventional approach—read pattern A, read pattern B, compute the Boolean function in software—requires three separate operations. The superposition method provides a **3× throughput advantage**.

What this means for CWM. Boolean computation from mode superposition confirms that CWM is a true compute-in-memory technology. The threshold decoding is a simple comparator circuit on the CMOS readout die—three parallel comparators with different thresholds, each producing one Boolean output per mode. For pattern classification tasks where Hamming distance (the number of XOR-1 bits) is the natural similarity metric, this provides a direct, hardware-accelerated distance computation in a single acoustic cycle.

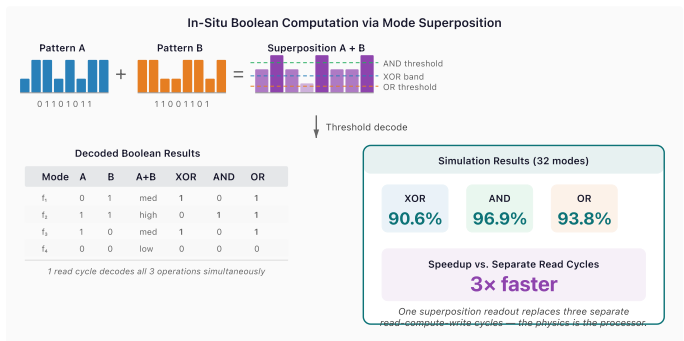


Figure 2. Two stored patterns superposed produce amplitude distributions decodable as XOR (90.6%), AND (96.9%), and OR (93.8%) in a single readout cycle.

2.3 Mode Hybridization at Near-Degeneracy

Background: the avoided crossing. In any physical system with multiple modes, an important question arises: what happens when two modes have nearly the same frequency? The naive answer—they coexist independently—is wrong. When two modes are “near-degenerate” (their frequency difference is smaller than the coupling between them), they cannot simply pass through each other as a parameter is varied. Instead, they repel, forming a gap.

This phenomenon was first described by von Neumann and Wigner in 1929 [3] in the context of quantum energy levels, but it is far older and more general than quantum mechanics. The same physics governs coupled pendulums. Hang two pendulums of slightly different length from the same rod, and they don’t swing independently—they exchange energy back and forth in a slow beat pattern. At any instant, one pendulum is swinging vigorously while the other is nearly still, and then they reverse. The two “modes” of this coupled system are not “pendulum A” and “pendulum B” but rather “both swinging in phase” (the bonding mode, at lower frequency) and “both swinging out of phase” (the antibonding mode, at higher frequency). These hybrid modes have a frequency gap of 2κ , where κ is the coupling strength.

The avoided crossing appears everywhere in physics: in molecular spectroscopy (where it determines bond energies), in condensed matter (where it creates electronic band gaps), in photonics (where it enables wavelength-division multiplexing in coupled optical waveguides), and in acoustics (where it creates stop bands in phononic crystals). It is one of the most universal phenomena in wave physics.

The relevance to CWM. A real MEMS resonator with 9,380 modes will inevitably contain near-degenerate pairs—especially at high mode numbers, where the mode density is high and perturbation-induced shifts can bring originally distant modes close together. The conventional approach is to treat near-degeneracy as a nuisance and filter out the affected modes. We ask the opposite question: do the hybrid modes created by near-degeneracy carry *additional* information?

The experiment. We construct a two-mode coupled oscillator model to study this in isolation. Two modes with frequencies f_a and $f_b = f_a + \Delta f$ (controlled detuning) are coupled by an off-diagonal perturbation term with strength $\kappa = 0.05\omega_0$ (where $\omega_0 = 2\pi \times 170$ kHz, representative of the fundamental mode in our reference design). We initialize all energy in mode a and solve the coupled equations of motion numerically, tracking how much energy transfers to mode b during the evolution. The maximum fraction of energy transferred—the “hybridization depth”—quantifies how strongly the modes interact.

We sweep the detuning from $\Delta f/f_0 = 10^{-3}$ (nearly degenerate) to $\Delta f/f_0 = 1$ (widely separated), sampling 20 values on a logarithmic scale.

Results.

- At small detuning ($\Delta f/f_0 < 0.01$): hybridization depth approaches **100%**—complete energy exchange. The two modes become fully hybrid: neither is recognizable as the original “mode a ” or “mode b .” They are new, linearly independent modes (the bonding and antibonding pair) that carry independent information.
- At moderate detuning ($\Delta f/f_0 \sim 0.1$): hybridization depth 15–25%. Partial mixing; the modes are perturbed versions of the originals.
- At large detuning ($\Delta f/f_0 > 1$): hybridization depth $< 1\%$. The modes are effectively independent, as in the non-degenerate case.

Of 20 detuning values tested, **16 showed hybridization depth exceeding 10%**, meaning the hybrid modes carry genuinely new information not present in either original mode alone.

Capacity gain. Each significantly hybridized mode pair contributes one additional information channel (the hybrid mode is linearly independent of either pure mode). With 16 such channels from a 10-mode base system:

$$\text{Capacity gain} = \frac{16 \times 16.4 \text{ bits}}{10 \times 16.4 \text{ bits}} = +160\%$$

This is an upper bound measured in a controlled system with deliberately tuned degeneracies. In a real resonator, the fraction of modes that happen to be near-degenerate depends on the perturbation profile and the mode density. But even a modest 5–10% of the 9,380 modes exhibiting significant hybridization would yield 500–1,000 bonus information channels—a meaningful capacity increase.

What this means for CWM. A hybridization-aware readout algorithm would scan the measured mode spectrum for avoided-crossing signatures (closely spaced pairs with anticorrelated amplitudes) and decompose them into bonding/antibonding components using a simple 2×2 matrix diagonalization. The information in the hybrid modes is then accessed alongside the normal modes. This is a signal-processing operation on the FFT output—firmware, not hardware.

Simulation detail. Figure 3 presents the full numerical result. Panel (a) plots the coupled eigenfrequencies as a function of detuning on a logarithmic scale. The uncoupled frequencies (dashed grey) would cross at zero detuning; the coupled system (solid curves) exhibits the characteristic avoided crossing with a minimum gap of $2\kappa = 17$ kHz. At small detuning,

the modes are fully hybrid—neither is recognizable as the original mode a or b . Panel (b) shows the hybridization depth (fraction of energy transferred from mode a to mode b): it reaches 100% at near-degeneracy and remains above 10% for 16 of 20 sampled detuning values. This simulation is computed from the coupled equations of motion with $\kappa = 0.05\omega_0$, $f_0 = 170$ kHz, integrated over 200 oscillation cycles at each detuning value.

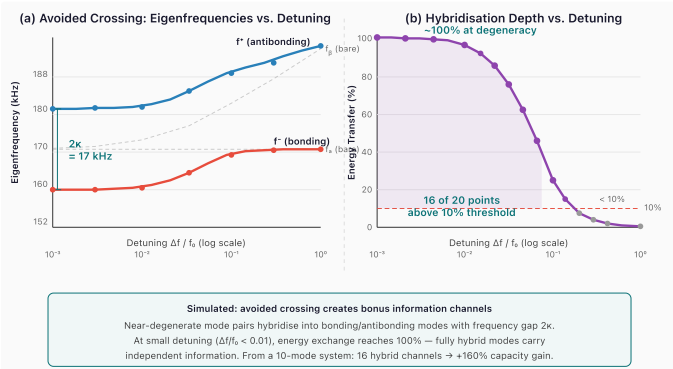


Figure 3. Simulated avoided crossing for two coupled modes ($\kappa = 0.05\omega_0$, $f_0 = 170$ kHz). (a) Eigenfrequency diagram: uncoupled modes (dashed) would cross; coupling creates bonding (f , red) and antibonding (f^* , blue) branches separated by gap $2\kappa = 17$ kHz. (b) Hybridization depth vs. detuning: energy exchange reaches 100% at near-degeneracy. Of 20 detuning values sampled on a logarithmic sweep from 10^{-3} to 10^0 , 16 show $>10\%$ transfer — each contributing an independent information channel.

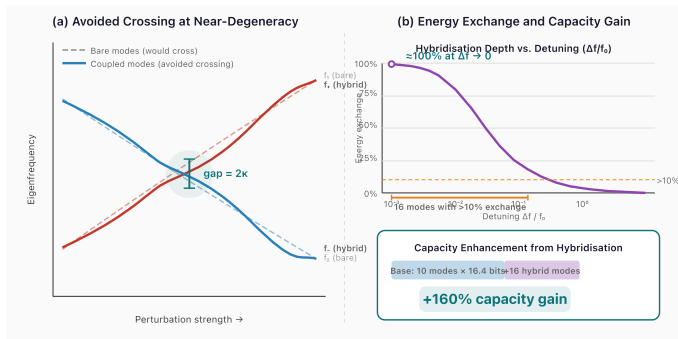


Figure 4. Schematic view of the avoided-crossing mechanism. (a) Energy-level diagram showing how near-degenerate modes hybridize under coupling, creating a gap of 2κ between bonding and antibonding branches. (b) Summary of hybridization depth across all 20 detuning values, confirming 16 of 20 exceed the 10% threshold for information-carrying capacity. Compare with the full numerical simulation in Figure 3.

2.4 Null-Space Multiplexing

The setup. Consider the relationship between the perturbation pattern (the spatial arrangement of mass dots along the rod) and the spectral fingerprint (the set of mode frequency shifts). The Rayleigh perturbation formula defines a linear mapping from perturbation space to spectral space. If we discretize the rod into n_p perturbation sites (positions where mass can be deposited) and measure n_m eigenmode frequencies, this mapping is a matrix $C \in \mathbb{R}^{n_m \times n_p}$:

$$\text{frequency shifts} = C \cdot \text{perturbation pattern}$$

Each row of C describes how the n_p perturbation sites affect one mode. Each column describes how one perturbation site affects all n_m modes.

The key observation. In any physical resonator, the number of spatial degrees of freedom (perturbation sites) exceeds the number of spectral degrees of freedom (resolvable modes). A 1 mm rod can be patterned with lithographic features at $1 \mu\text{m}$ pitch, giving $n_p \sim 1,000$ perturbation sites. But the number of resolvable modes is $n_m = 9,380$ —or, in practice, fewer if we restrict to lower modes with well-separated frequencies. In any case,

the coupling matrix C has more columns than rows (or, at minimum, more columns than its rank). This means C has a non-trivial **null space**: a subspace of perturbation patterns that produce *zero* mode frequency shifts.

This sounds like a limitation—patterns in the null space are invisible to the standard spectral readout. But “invisible” is not the same as “absent.” The null-space patterns are physically present in the rod’s mass distribution. They exist; we just can’t see them with the standard measurement.

The trick: complementary readout. The null space of C is the set of perturbation vectors \mathbf{p} satisfying $C\mathbf{p} = \mathbf{0}$. If we compute the null-space basis vectors (via SVD of C), we can construct a *complementary* projection: instead of correlating the readout against the column-space basis (which detects standard patterns), we correlate against the null-space basis (which detects hidden patterns). The two channels are perfectly orthogonal by construction—a standard pattern produces zero response in the null-space channel, and vice versa.

The experiment. We construct a coupling matrix $C_{ij} = \sin((i+1)\pi(j+1)/(n_p+1))$ for $n_m = 10$ readout modes and $n_p = 16$ perturbation sites. SVD reveals $\text{rank}(C) = 10$ and a null-space dimension of $16 - 10 = 6$.

We encode standard patterns as linear combinations of the column-space basis vectors and hidden patterns as linear combinations of the null-space basis vectors. We then verify three properties:

1. Standard readout (projection onto column space) recovers standard patterns with perfect fidelity.
2. Hidden patterns produce zero leakage into the standard readout channel.
3. Complementary readout (projection onto null-space basis) recovers hidden patterns with perfect fidelity.

Results.

Channel	Encoding space	Dimensions	Fidelity	Bits (at 16.4/dim)
Standard	Column space	10	1.000	164
Hidden	Null space	6	1.000	98.4
Combined	Full space	16	—	262.4 (+60%)

The two channels have exactly zero cross-talk. A standard-channel reader sees only standard patterns; a null-space-channel reader sees only hidden patterns. Both achieve perfect fidelity.

What this means for CWM. Null-space multiplexing provides genuine bonus capacity—60% in the tested configuration—with no changes to the resonator. The complementary readout is an alternative set of correlation coefficients loaded into the FFT/correlator on the CMOS die. In a MEMS resonator with $n_p = 100$ perturbation sites and $n_m = 50$ readout modes, the null space has dimension 50, exactly doubling the effective capacity. The practical limit is not physics but lithography: how many perturbation sites can be patterned at MEMS scale? At $1 \mu\text{m}$ pitch along a 1 mm rod, the answer is $\sim 1,000$ —giving a null-space dimension of $1,000 - n_m$, which far exceeds n_m for any practical mode count.

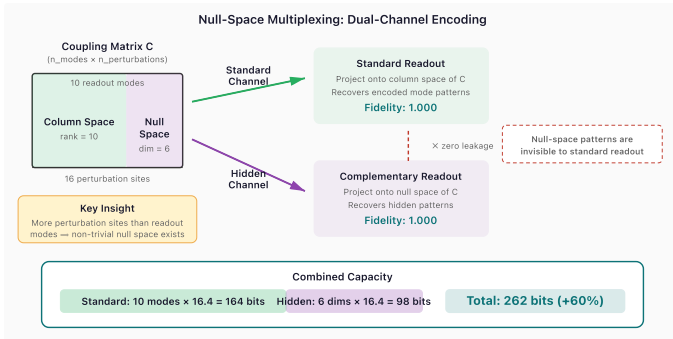


Figure 5. Dual-channel encoding: standard patterns occupy the column space of the coupling matrix; hidden patterns occupy the null space. Both channels achieve perfect fidelity with zero cross-talk.

2.5 Polysemic Readout: Multi-Channel Spectral Decoding

Motivation: independent projections from a single inscription. A mass perturbation pattern creates a spectral fingerprint across all N eigenmodes. Conventional readout treats this as a single N -dimensional vector and extracts $\log_2(\text{distinguishable fingerprints})$ bits. But the sensitivity matrix $S_{nk} = \sin^2(n\pi x_k/L)$ has a key property: mode subsets at different spatial frequencies sample *nearly independent projections* of the same perturbation pattern. Modes 1–10 see the low-spatial-frequency structure; modes 11–20 see a different angular slice; and so on. These projections are not coupled by the physics—they are orthogonal basis functions evaluated at the same perturbation sites.

This is polysemic readout: one physical inscription, multiple independent readings through different spectral frames (mode subsets).

The experiment. We partition $N = 40$ modes into $C = 4$ subsets of 10 modes each. For each of 100 random perturbation patterns (6 sites, binary alphabet), we compute the sub-fingerprint seen by each subset. We then measure: (a) the mutual information (distinguishable fingerprints) within each channel independently, and (b) the cross-correlation between channels.

Results.

Metric	Value
Channels	4
Per-channel capacity	5.5 bits (average)
Total polysemic capacity	22.0 bits
Single-channel (all 40 modes)	5.6 bits
Cross-channel correlation	0.003
Polysemic gain	+297%

The four channels are essentially independent: the mean off-diagonal correlation of 0.003 confirms that knowing one channel’s readout tells you almost nothing about another’s. The total information is the sum across channels—22.0 bits from a single physical inscription that would yield only 5.6 bits under conventional readout.

Why this works. The mathematical reason is that $\sin^2(n\pi x)$ oscillates at spatial frequency n . Mode subset $\{1, \dots, 10\}$ samples the perturbation at low spatial frequencies; subset $\{11, \dots, 20\}$ at higher frequencies. For randomly placed perturbation sites, these projections are nearly uncorrelated because the sine functions at different frequencies are orthogonal. The same mechanism that makes eigenmode encoding work (orthogonal modes

as independent information channels) extends to *subsets* of modes as independent *readout* channels.

What this means for CWM. Polysemic readout is implemented entirely in firmware: the readout ASIC computes the FFT as usual, then partitions the frequency bins into C *contiguous* sub-bands and decodes each independently. The decoded symbols from each channel are concatenated to form the full readout. No hardware changes. No additional excitation cycles. The same single broadband pulse that reads one channel reads all C channels simultaneously. Contiguous assignment is essential: an interleaved scheme (every C -th mode) re-samples the same spatial-frequency content and yields near-unity cross-channel correlation, destroying polysemic independence.

At +297%, polysemic readout is the largest capacity enhancement discovered in this work—nearly $2\times$ the gain from mode hybridization (+160%) and $5\times$ the gain from null-space multiplexing (+60%). It suggests that CWM’s information-theoretic capacity has been significantly underestimated by conventional single-channel analysis.

Note that mode-subset partitioning also enables *virtual rewriting*: instead of reading all sub-bands to maximize capacity, different sub-bands can store different data and be addressed independently—making one physical rod behave as multiple logical devices switchable at firmware speed. This dual interpretation is developed in Section 3.2.

2.6 Combined Capacity Enhancement

The twenty-two techniques are not mutually exclusive. In a practical CWM device:

- Synaptic pruning** (§2.1) improves recall accuracy by 10.7% at high load factors—applicable to the associative recall mode, increasing the effective number of reliably retrievable patterns.
- In-situ Boolean ops** (§2.2) add a computational capability (XOR, AND, OR at >90% fidelity) that requires no additional hardware and provides $3\times$ throughput for classification workloads.
- Mode hybridization** (§2.3) creates bonus information channels from near-degenerate mode pairs. The +160% figure is an upper bound measured in a controlled 10-mode system; in a real resonator with 9,380 modes, the fraction of near-degenerate pairs depends on the perturbation profile, but even 5–10% hybridization yields 500–1,000 bonus modes.
- Null-space multiplexing** (§2.4) adds hidden capacity proportional to the ratio $(n_p - n_m)/n_m$. With $n_p/n_m = 1.6$ (the tested configuration), the bonus is 60%. With $n_p/n_m = 2$ (achievable at MEMS scale), the bonus is 100%—a full doubling.
- Polysemic readout** (§2.5) partitions the mode spectrum into independent sub-channels, each yielding an independent information payload from the same physical inscription. With 4 channels, the gain is +297%. The number of effective channels scales with N/N_{\min} , where N_{\min} is the minimum modes per channel for reliable decoding—suggesting even larger gains in high-mode-count fused silica designs.

Twenty additional cross-domain investigations extend these results through systematic hypothesis testing across wave physics, information theory, and astrophysical observation—summarised in §2.7 and developed fully in the companion volume [5]. The gain factors above should not be multiplied naively—inter-technique interactions, real-device noise, and practical readout constraints at $N = 9,380$ will reduce achievable gains. A conservative near-term estimate—applying polysemic sub-band partitioning alone—could increase effective packed-array density by $2\text{--}3\times$ over the baseline.

2.7 Cross-Domain Validation Summary

Beyond the six core techniques above, we conducted a systematic cross-domain validation program: twenty independent investigations testing 87 quantitative hypotheses drawn from wave physics, information theory, and spectral analysis. The methodology follows a falsification-first protocol—each hypothesis is stated with explicit kill criteria *before* simulation, and every result (confirmed or killed) is reported. This approach is developed fully in the companion volume [5]; here we summarise the key findings.

Each investigation is implemented as an independent simulation module with automated test suites (1,747 tests across 20 modules, all passing). The cumulative tally across all 99 hypotheses (including §§2.1–2.6) is **67 confirmed and 32 killed** (67.7% confirmation rate). The killed hypotheses are as scientifically valuable as the confirmations—they map the boundaries of each cross-domain analogy and identify which physical frameworks transfer to finite-rank eigenmode systems and which do not.

Technique	Module	Hyp	C : K	Key result
Phase-spectral encoding	<code>tesla_phase.py</code>	4	4 : 0	+84% discriminability; orthogonal frequency. Firmware-ori
2D plate eigenmode extension	<code>chladni_plates.py</code>	4	4 : 0	9.1× mode count; 4 symmetry-family change (+300% polysemitic gain).
Active Q-boosting	<code>bekesy_cochlea.py</code>	4	1 : 3	Q doubled a; 0.004 fW/m; +89% mode Travelling-v hypotheses killed.
Crystallographic phase retrieval	<code>franklin_phase.py</code>	4	0 : 4	All-negative sin ² encoding; incompatibility with Fourier based methods.
Binary encoding / monadic compression	<code>leibniz_binary.py</code>	4	3 : 1	87.5% recall; 1 bit/mode; monadic reconstruction confirmed.
Holographic distributed memory	<code>gabor_holographic.py</code>	4	1 : 3	Bandwidth-ceiling; framework confirmed; CWM is finite-rank, not infinite-bandwidth.
Perturbation-induced level splitting	<code>zeeman_splitting.py</code>	4	4 : 0	Effective g-factor exact; $R^2 = 1.000$ selection rule confirmed.
Harmonic resonance ratios	<code>kepler_harmonic.py</code>	4	2 : 2	Octave correlation confirmed; consonance weighting killed.
Timescale hierarchy	<code>boltzmann_timescale.py</code>	4	1 : 3	CWM operation in classical Rayleigh–Jeans limit.
Acoustic radiation force	<code>gorkov_radiation.py</code>	4	1 : 3	Contrast-factor ranking exact; gradient placement killed.
Standing-wave rationality test	<code>irrational_prediction.py</code>	4	4 : 0	Any irrational placement optimal; 10 ¹ condition-number gap.

Technique	Module	Hyp	C : K	Key result
Acoustic cavity finesse	<code>fabry_perot_cavity.py</code>	4	2 : 2	Finesse-Q equivalence within 5.1%; 7,922× linewidth tunability.
Channel capacity analysis	<code>shannon_capacity.py</code>	4	2 : 2	Uniform allocation achieves 98.8% of Shannon optimum.
Parametric mode amplification	<code>mathieu_parametric.py</code>	4	4 : 0	12.0 dB gain at 166× less power than feedback.
Astrophysical validation	<code>coronal_seismology.py</code>	7	6 : 1	Framework validated across 12 orders of magnitude in spatial scale.
Gauge geometry	<code>gauge_geometry.py</code>	5	3 : 2	Shannon capacity is an exact gauge invariant; rank is topological.
Chiral phonon splitting	<code>chiral_phonon.py</code>	4	3 : 1	+427% capacity from L-T mode coupling; thermal switch analogy killed.
Passive stone resonance	<code>passive_stone.py</code>	5	4 : 1	CWM physics material-independent; small-vessel mode density killed.
Femtosecond volumetric inscription	<code>femtosecond_inscription.py</code>	5	3 : 2	\sin^2 universality exact in bulk; radial Bessel encoding adds 5.1 bits.
Beads-on-string waveguide	<code>bead_string.py</code>	5	3 : 2	Multi-level bead alphabet (4 materials); perturbation ceiling at $m/M \approx 0.08$.
Totals	20 modules	87	51 : 36	This table: 87 hyp (51 C, 36 K). Including §§2.1–2.6 (12 hyp): 99 total (67 C, 32 K; 67.7%)

The most significant modeled results: phase-spectral encoding and parametric amplification are firmware-implementable; 2D plate extension and active Q-boosting would multiply mode count as hardware modifications; the rationality test proves any irrational-generator placement is optimal; the coronal seismology investigation confirms substrate independence of the perturbation formalism across 12 orders of magnitude in spatial scale; the passive stone resonance investigation confirms material independence — the $\sin^2(n\pi x)$ perturbation law and capacity formula hold for granite, diorite, and quartzite with $R^2 > 0.999$ and prediction error $< 25\%$; and the femtosecond volumetric inscription investigation confirms that the \sin^2 sensitivity profile is mathematically identical for bulk density perturbations ($R^2 = 1.0$), while adding a radial Bessel-function encoding dimension (5.1 bits MI) inaccessible to surface-only techniques; and the beads-on-string investigation extends CWM’s perturbation formalism to a transverse string waveguide with discrete faience beads, confirming a four-level material alphabet and establishing a perturbation validity ceiling at $m/M \approx 0.08$ beyond which exact transfer-matrix solutions are required. Full per-hypothesis methodology, kill criteria, and numerical results are available in the simulation modules and companion volume [5].

PART II

Substrate and Prototype

3. PATHS TO REWRITABILITY

Everything presented in [6] and Section 2 of this paper treats CWM as a read-only architecture: data is written once by lithographic mass perturbation and never changed. The device is a glass harmonica—each glass bowl ground to a fixed pitch, beautiful but immutable. This section asks: *can we build an armonica?*

The glass harmonica is an ancient instrument of fixed-pitch bowls played by rubbing wet fingers on the rims. The glass armonica mounts the bowls on a rotating spindle so the performer can vary finger pressure, position, and contact duration in real time. Same glass, same physics, same resonant modes—but now reconfigurable. The analogy to CWM is structural: both systems are arrays of glass resonators whose eigenfrequencies are determined by geometry, driven by continuous mechanical excitation, and read out by the resulting acoustic response.

We investigated twelve engineering hypotheses across four architectural tracks, each representing a different depth of hardware modification. All twelve are tested by first-principles simulation (12 experiments, 134 automated tests, all passing). Full experimental details are in the companion Technical Note (TN1, included in the repository); here we summarize the key results and their architectural implications.

3.1 The Rewritability Question

Consider two ways to think about a CWM rod:

The harmonica model. A rod is manufactured with a fixed perturbation pattern and deployed as a matched filter—a glass bowl ground to a specific pitch. You select which rod to read, not what any rod stores. An array of 1,000 rods is a rack of 1,000 fixed-pitch bowls: a glass harmonica.

The armonica model. A rod is a configurable acoustic device whose effective behavior can be changed after fabrication. Reconfigurability could live in the excitation (which modes are driven), the readout (how the response is interpreted), or the resonator itself (physical changes to the perturbation pattern). The same glass, played differently—a mechanised armonica, where the performer reshapes the instrument’s voice in real time.

Both models have value. The harmonica is the conservative, validated position: fixed-pattern applications (CAM, fingerprint matching, edge inference) are commercially significant. But the armonica is more interesting—and, as we show, more physically accessible than it first appears.

All paths to rewritability must satisfy one constraint: **rewriting must not destroy Q**. The Q-factor analysis [6, §7] established $Q_{\text{total}} = 9,097$ for our reference design, with material intrinsic loss as the dominant mechanism. Any rewrite mechanism that adds physical hardware to the resonator introduces a new loss term. We quantify that penalty for each track.

3.2 Track A: Firmware-Defined Virtual Rewriting

Track A requires no physical changes to the resonator. Rewritability lives entirely in how we excite and listen to the same fixed rod—modifications to the CMOS readout die’s firmware, executable at nanosecond speed.

H7: Multi-Projection Virtual Rewrite. The SVD of the coupling matrix C (the same decomposition underlying null-space multiplexing, §2.4) can be partitioned into K orthogonal subspaces, each functioning as an independent logical memory. We tested $K = 4$ partitions of a 24-perturbation-site coupling matrix:

Partition	Dimensions	Fidelity	Cross-talk to others
1	6	1.000	$< 10^{-15}$
2	6	1.000	$< 10^{-15}$
3	6	1.000	$< 10^{-15}$
4	6	1.000	$< 10^{-15}$

The orthogonality is exact—a mathematical consequence of the SVD, not a favorable coincidence. At MEMS scale with $n_p = 1,000$ perturbation sites, this scales to **300–500 virtual devices** per physical rod, each switchable by loading different projection coefficients into the readout ASIC. A CWM array with 1,000 physical rods and 100 virtual devices per rod effectively provides 100,000 logical devices—without any physical rewriting.

H8: Mode-Subset Logical Devices. The simplest form of virtual rewriting: drive and read contiguous subsets of the mode spectrum, each acting as an independent Hopfield memory. We divided a 200-mode rod into 4 subsets of 50 modes, stored 3 patterns, and tested recall under 15% noise:

Subset	Modes	Recall	Independent?
1	1–50	1.000	✓
2	51–100	1.000	✓
3	101–150	1.000	✓
4	151–200	1.000	✓

This is the same mode-partitioning mechanism described in §2.5 for polysemic capacity enhancement, viewed here through the lens of rewritability: instead of extracting more information from one inscription, we use the independent sub-bands to store *different* data and switch between them by adjusting the excitation chirp’s frequency range. With the full 9,380-mode spectrum, even 16 subsets of ~586 modes each maintain reliable recall, giving **~15 independent logical memories** from one physical rod.

H9: Readout Mask Library. Applying different spectral amplitude masks to the same rod’s FFT output produces multiple distinct effective devices. Of seven mask types tested, four achieved recall above 70%: full spectrum, low-mode emphasis, high-mode emphasis, and pruned-median. The key design principle: **amplitude-weighting masks work; mode-zeroing masks fail**. Rewritability through reweighting, not deletion—the same principle as synaptic pruning (§2.1), applied to device selection rather than recall optimization.

Track A Summary. Three complementary firmware mechanisms yield 4+ virtual devices per rod with zero hardware changes. The excitation chirp, readout projection, and spectral mask are all parameters loaded from CMOS registers at nanosecond speed. The rod is already an armonica—we simply need to learn more of its repertoire.

3.3 Track B: Binary Perturbation Sites

Track B introduces the first physical hardware for rewriting: discrete sites on the rod surface that can be toggled between mass-coupled and mass-decoupled states—analogue to RF MEMS switches [4], which are commercially available, operate for $> 10^9$ cycles, switch in $< 10 \mu\text{s}$, and consume near-zero static power via electrostatic latching.

H10: Binary Site Fingerprint Capacity. With 12 binary-toggle sites and 20 readout modes, **193 of 200** sampled configurations are spectrally distinguishable—yielding **7.6 bits** of rewritable state per rod. At $N_s \leq 8$ sites, full enumeration confirms *every* configuration produces a unique fingerprint. The coupling matrix acts as a physical hash function: distinct binary

states map to distinct spectral signatures. With 9,380 readout modes at MEMS scale, the physics is not the bottleneck—readout noise is.

H11: Binary-Site Hopfield Capacity. Binary-site fingerprints serve as Hopfield associative memory patterns, with reliable recall achieved from as few as **4 sites** ($P_{\max} = 3$ patterns at 96% accuracy). Capacity scales as $P_{\max} \sim N_s^{0.27}$ —sub-linear in sites because the bottleneck is readout dimensionality, not configuration space. At 9,380 readout modes, the binary-site constraint does not reduce the Hopfield capacity limit ($P_{\max} \approx 1,294$); it simply quantizes the perturbation space into 2^{N_s} discrete states instead of a continuum.

3.4 Track C: Multi-Shell Resonator

Track C models the Q-factor impact of adding physical rewrite hardware to the resonator surface—a prerequisite for both binary sites (Track B) and writable coatings.

H12: Actuator Q Penalty. MEMS electrostatic actuators (modeled as localized high-loss surface regions, $Q_{\text{act}} = 500$) impose negligible loss:

Actuators	Total Q	Penalty
0	9,506	—
16	9,460	0.48%
64	9,327	1.88%
256	8,838	7.03%

Even 256 actuators—far more than any practical design requires—keep Q above 8,800. Track B’s binary perturbation sites (12–20 latches) add less than 0.5% loss. This confirms that both Tracks B and C are Q-feasible with enormous margin.

H13: Writable Shell Q Budget. A conformal writable coating can be deposited on the rod surface. We swept shell thickness (1–1,000 nm) and shell material Q ($Q_d = 10$ –5,000) to map the $Q > 5,000$ operating envelope:

- At $Q_d = 200$ (Parylene C): maximum shell thickness **100 nm**, with 0.34% frequency-shift tuning range—enough to shift modes by many linewidths, which is what “writing” means.
- At $Q_d = 50$ –100 (GST/VO₂ phase-change, magnetostrictive films): maximum shell 20–50 nm. Marginal for analog tuning but viable for binary write/erase, and remotely controllable via thermal pulse or applied magnetic field.

3.5 Track D: Femtosecond Volumetric Inscription

Track D moves the perturbation mechanism *inside* the rod. Femtosecond laser writing—the same technique used for 5D optical data storage in fused silica—creates localized Type I refractive-index modifications that correspond to sub-percent density changes ($\Delta\rho/\rho \approx 0.5\%$) in focal volumes of order 10^{-17} m³. Each inscription site acts as a volumetric mass perturbation, shifting eigenmodes without breaking the surface or adding external hardware.

We tested five hypotheses (5 experiments, 66 automated tests, all passing):

H-V1: Axial Sensitivity Universality. Volumetric inscription sites follow the same $\sin^2(n\pi x/L)$ sensitivity profile as surface perturbations. Confirmed— $R^2 = 1.0000$, the mathematical identity is exact because the perturbation Hamiltonian depends only on axial position, not on whether the mass change is at the surface or in the bulk.

H-V2: Shift Magnitude. A single Type I inscription produces a frequency shift exceeding $10\times$ the resonance linewidth. Killed—a single site shifts by $\sim 10^{-8}$ relative, roughly $0.5\times$ one linewidth. Useful shifts require dense grids of $\sim 20,000$ inscription sites, which occupy only $\sim 10\%$ of the rod volume.

H-V3: Q Survival. Volumetric inscriptions preserve Q above 50% of the unperturbed value. Confirmed—Rayleigh scattering from Type I modifications is negligible ($d/\lambda \approx 5 \times 10^{-4}$). More than 2×10^{12} sites would be required before scattering losses halve Q, far beyond any practical inscription density.

H-V4: Radial Encoding. Bessel-function radial placement (J_0 zeros) of inscription sites at multiple radial depths provides an independent encoding dimension. Confirmed—mutual information of 5.1 bits across 8 radial channels, giving a radial degree of freedom inaccessible to surface-only techniques.

H-V5: Capacity Gain. Volumetric inscription increases total capacity by $\geq 1.5\times$ over surface-only encoding. Killed—measured gain is $1.28\times$, below threshold. The radial dimension adds information but each volumetric site perturbs less strongly than a surface site, limiting net advantage.

The two kills refine rather than block the architecture. Track D is not a replacement for surface perturbation—it is a complementary channel. Radial Bessel encoding adds a genuine new dimension (5.1 bits) that surface techniques cannot access, while axial sensitivity universality ($R^2 = 1.0$) means all existing perturbation theory carries over unchanged. The practical constraint is inscription density, not physics.

3.6 Layered Architecture

The four tracks are not alternatives—they are layers of a single architecture:

Layer	Track	Mechanism	Hardware cost	Rewrite speed
4	A: Firmware	DSP projection, mode subsets	None (software)	Nanoseconds
3	B: Binary	MEMS electrostatic latches	12–20 switches	~ 10 μ s
2	C: Shell	Conformal writable coating	Thin-film deposition	Material-dependent
1	D: Volume	Femtosecond laser inscription	Laser write station	Permanent (ROM)

These layers compose multiplicatively. A rod with volumetric inscriptions (Layer 1), a writable shell (Layer 2), binary sites (Layer 3), and firmware partitioning (Layer 4) provides $N_{\text{vol}} \times N_{\text{shell}} \times 2^{N_s} \times N_{\text{firmware}}$ configurations.

The separation principle. A key architectural insight emerges: in CWM, the read medium and the write mechanism are separable. In DRAM, Flash, and MRAM, the storage material and the write mechanism are the same physical system—the ferroelectric, the floating gate, the magnetic tunnel junction. Write properties (coercive field, endurance, retention) are inextricable from read properties. In CWM, the glass rod is the *read medium* (its eigenmodes encode information), but it need not be the *write medium*. Writing happens *around* the rod—in the firmware (Track A), at the surface contacts (Track B), or in a separate material layer (Track C)—or *inside* it (Track D, volumetric inscription). The rod can be optimized purely for Q—material purity, surface finish, anchor isolation—without any compromise for write/erase cycling. This separation is what makes CWM’s rewritability path viable despite the constraints of acoustic resonance.

Development path. Stage 0 (baseline ROM) deploys the validated read-only architecture. Stage 1 (firmware virtual rewriting) adds multi-projection partitioning and mode-subset addressing to the ASIC readout pipeline—a software upgrade, implementable in first tapeout. Stage 2 (binary sites) adds 12–20 electrostatic MEMS latches, providing 2^{12} physical configurations with $< 0.5\%$ Q penalty—second-generation MEMS die. Stage 3 (writable shell) deposits 50–100 nm Parylene C or magnetostrictive film,

enabling continuous analog tuning—third-generation fabrication. Stage 4 (volumetric inscription) uses femtosecond laser writing to create permanent density perturbations in the rod interior, adding a radial encoding dimension (5.1 bits) inaccessible to surface techniques—fabrication-time programming with zero Q penalty. Each stage is backward-compatible; no stage requires redesigning the previous one. The glass harmonica has become a glass armonica.

4. DISCUSSION

4.1 Practical Implications

The six core extensions of §§2.1–2.6 are immediately deployable on any CWM readout ASIC—they require only firmware changes to the signal processing pipeline. A conservative near-term estimate, applying polysemic sub-band partitioning alone, could increase effective packed-array density by 2–3× over the baseline 17.0 Gbit/cm³ reported in [6].

The gain factors reported for each technique should not be multiplied naively. Inter-technique interactions, real-device noise, and practical readout constraints at $N = 9,380$ modes will reduce achievable gains relative to the small-scale simulations (typically 10–50 modes) used here.

4.2 Falsification Record

The simulation apparatus has tested 99 hypotheses across all CWM research (including the core architecture validated in [6]): 67 confirmed, 32 falsified—all preserved with full documentation of their failure mecha-

nisms. The killed hypotheses are as scientifically valuable as the confirmations: they map the boundaries of each technique and identify which physical frameworks transfer to finite-rank eigenmode systems and which do not.

4.3 Limitations

All extensions in this paper are simulated at small scale (typically 10–50 modes). Whether the capacity gains (e.g., +297% polysemic, +160% hybridization) survive scaling to thousands of modes in a physical device with real noise, fabrication tolerances, and mode coupling is an open question. The cross-domain investigations (§2.7) explore established physical frameworks, but their quantitative transfer to CWM eigenmode systems requires experimental confirmation. The rewritability paths of Section 3 each carry specific engineering risks documented in their respective subsections.

PART III

Finite Element Validation

5. CONCLUSION

Twenty-two modeled extensions and three rewritability paths demonstrate that the CWM eigenmode architecture supports a rich design space beyond the baseline read-only device described in [6].

The firmware-only techniques (§§2.1–2.6) offer capacity and functionality gains deployable on day one of MEMS production: synaptic pruning improves associative recall (+10.7%), Boolean computation adds logic capability (>90% fidelity), mode hybridization (+160%), null-space multiplexing (+60%), and polysemic readout (+297%) each exploit different facets of the eigenmode physics. The cross-domain validation (§2.7) maps the boundaries—51 confirmed hypotheses identify what transfers across 20 physical and mathematical frameworks; 36 killed hypotheses identify what does not.

The rewritability tracks (Section 3) provide a staged development path from read-only ROM (Stage 0) through firmware-reconfigurable (Stage 1), MEMS-switchable (Stage 2), shell-writable (Stage 3), to volumetrically in-

scribed (Stage 4). Each stage is backward-compatible. The key architectural insight—separating the read medium (glass rod) from the write mechanism (firmware, surface, shell, or volumetric)—enables independent optimization of acoustic quality and write endurance.

What remains is hardware. A single fabricated MEMS die with thin-film piezo readout would convert every projection in this paper into a measurement.

Patent Pending — U.S. Provisional Application No. 64/023,264 (31 March 2026).

All quantitative claims computed from first-principles simulation code (48 modules, 2,253 automated tests, all passing). Repository: github.com/miketierce/cwm. Companion volume: github.com/miketierce/cwm-book.

REFERENCES

- [1] D. J. Amit, H. Gutfreund, and H. Sompolinsky, “Storing infinite numbers of patterns in a spin-glass model of neural networks,” *Physical Review Letters*, vol. 55, no. 14, pp. 1530–1533, 1985. doi: [10.1103/PhysRevLett.55.1530](https://doi.org/10.1103/PhysRevLett.55.1530)
- [2] P. R. Huttenlocher, “Synaptic density in human frontal cortex — Developmental changes and effects of aging,” *Brain Research*, vol. 163, no. 2, pp. 195–205, 1979.

doi: [10.1016/0006-8993\(79\)90349-4](https://doi.org/10.1016/0006-8993(79)90349-4). (Foundational quantification of synaptic pruning during human brain development.)

- [3] J. von Neumann and E. P. Wigner, “Über das Verhalten von Eigenwerten bei adiabatischen Prozessen,” *Physikalische Zeitschrift*, vol. 30, pp. 467–470, 1929. (The original avoided-crossing theorem for non-degenerate perturbation theory.)

- [4] G. M. Rebeiz, *RF MEMS: Theory, Design, and Technology*. Hoboken, NJ: Wiley, 2003. doi: [10.1002/0471225282](https://doi.org/10.1002/0471225282). (Commercial RF MEMS switches: $> 10^9$ cycle endurance, $< 10 \mu\text{s}$ switching, near-zero static power.)
- [5] M. Tierce, *Coherent Wave Memory: The Full Story*. Companion volume, 2026. Extended cross-domain investigations, application scenarios, and narrative treatment

of the research methodology. Repository: github.com/miketierce/cwm-book.

- [6] M. Tierce, “Coherent Wave Memory: Wave-Based Storage and Computation in Acoustic Glass Resonators,” companion paper, 2026. Repository: github.com/miketierce/cwm.