Coherent Wave Memory: Spectral Encoding, Classical Non-Separability, and Computation in Acoustic Glass Resonators

Mike Tierce

Independent Researcher

ORCID: 0009-0004-3869-958X | github.com/miketierce/cwm | Version 19 — June 2026

Abstract—
We demonstrate that fused silica plates driven at acoustic eigenfrequencies exhibit non-separable frequency×space correlations quantified by a CHSH-analog parameter $S = 2.827$ — 99.95% of the Tsirelson bound $2\sqrt{2}$ — at a significance exceeding 200,000 standard deviations above the separability limit. Five independent mode pairs all exceed $S = 2.82$, stable within 0.65% over 3.5 hours and 16.5 million drive cycles. This classical (intra-system) non-separability, following the framework of Qian & Eberly (2011) and Kagalwala et al. (2013), constitutes the first demonstration in an acoustic solid-state resonator and provides a measurable computational resource: different spatial receivers decode distinct information from the same broadband excitation, enabling spatial multiplexing at 100% classification accuracy.

Building on this physics, we report a complete experimental characterization of eigenmode-spectral encoding in a 38glassplateprototype :Q = 2{,}759 (intrinsic), 27resolvablemodesspanning30–120kHz, 100\sigma separation(80/80trials), and256 − pattern(8 − bit)multi − levelencodingwithzeroclassificationerror. Booleancomputation(AND/OR/X Q_\text{loaded} = 152–241, memorydepth\approx$ 2 steps) — a rate-limited engineering constraint, not a physics limitation.

A five-mechanism Q-factor model predicts $Q_{\text{total}} = 9{,}097$ for a 1 mm MEMS resonator, with material loss dominating (91% of budget). At MEMS rates (2 µs step interval vs. 52 ms decay time), the measured Lorentzian kernel and spatial transfer matrix would yield a functioning reservoir computer — a 2.5-million-fold improvement in the step/decay ratio over the current bench. All eigenfrequency predictions are independently validated by finite element analysis to 7 ppm.

No quantum nonlocality is claimed. All results are classical.

CLAIM-LABEL KEY

Measured = instrument reading on built hardware. **Derived** = mathematical consequence of measured quantities. **Projected** = engineering extrapolation to unbuilt MEMS. **Modeled** = simulation only, not yet on hardware.

PART I

Theory and Architecture

1. INTRODUCTION

1.1 The Memory Wall

In 1978, John Backus delivered his Turing Award lecture asking whether programming could be "liberated from the von Neumann style" [1]. His complaint was architectural: in the von Neumann model, every operation requires data to travel between physically separate storage and computation elements. The resulting bandwidth constraint — the "memory wall" identified by Wulf and McKee in 1995 [2] — has worsened with each process node. In modern data centers, more than half the energy budget is spent moving data, not computing on it.

Several in-memory computing architectures address this by co-locating storage and computation in resistive crossbar arrays [3]. These are genuine advances, but they still encode data as electrical states and manipulate those states with electrical signals.

1.2 Wave-Based Memory

We propose a fundamentally different physical encoding: store data as the eigenmode spectrum of a mechanical glass resonator, and compute by wave interference.

A glass plate vibrating at its eigenfrequencies supports N independent modes — each a separate frequency channel with its own spatial shape. Mass perturbations shift eigenfrequencies via the Rayleigh perturbation formula [7], creating unique spectral fingerprints. Reading requires only a broadband excitation and an FFT. Searching an array of resonators requires driving all elements simultaneously with a query pattern: the element whose stored fingerprint best matches the query resonates most strongly — a nearest-neighbor computation performed by wave physics in a single acoustic propagation cycle, with no processor, no bus, no software.

This paper reports the experimental validation of this architecture on a fused silica plate. The central discovery is that acoustic eigenmodes in glass plates exhibit non-separable frequency×space correlations ($S = 2.827$, five pairs, $>$200,000σ) that are computationally useful — enabling spatial multiplexing, content-addressable memory, and physical unclonable functions. We also identify the architecture's current limitations (temporal memory, rate constraints) and present the scaling argument to MEMS.

1.3 Summary of Results

Density definitions. *Active density:* bits ÷ resonator volume. *Packed-array density:* bits ÷ array volume at 2× diameter pitch. *Package density:* not estimated (requires fabricated device).

Parameter	Value	Label
Plate Q-factor (intrinsic)	2,759 at 35,840 Hz	Measured
Plate Q-factor (loaded)	152–241	Measured
Mode census	27 modes, 30–120 kHz	Measured
Single-capture discrimination	100% (80/80 trials, 193σ)	Measured
Multi-level encoding	8 levels × 4 modes = 256 patterns (8 bits), zero error	Measured
Boolean compute (AND/OR/XOR)	100% at 4 bits (16 patterns)	Measured
CHSH non-separability witness	$S = 2.827$ (best pair); 5/5 pairs > 2.82	Measured
Temporal stability (CHSH)	$S = 2.826 \pm 0.0003$ over 3.5 h	Measured
Endurance	16.5M cycles, 0.22% max drift	Measured
Bench SNR	42–55 dB (plate, Pico NCO drive)	Measured
Thermodynamic SNR ceiling	98.5 dB (derived from $\frac{1}{2}k_{\text{eff}}A^2/k_BT$)	Derived
Thermally stable modes (borosilicate)	9,380	Derived
Q_{total} (MEMS, 1 mm)	9,097	Modeled
Packed-array density (1 mm boro.)	17.0 Gbit/cm³	Projected
Packed-array density (0.5 mm SiO_2)	1.4 Tbit/cm³	Projected
Write energy	15 fJ/bit (physics-layer)	Projected
NARMA-10 reservoir (bench)	FAIL (NRMSE > 1.0)	Measured
NARMA-10 reservoir (simulated MEMS)	NRMSE = 0.39	Modeled

Validated vs. projected at a glance. The plate prototype, Q measurement, mode census, CHSH non-separability, Boolean compute, multi-level encoding, and endurance test are *measured*. Mode count formula and Shannon capacity are *derived*. All MEMS density, energy, latency, and Q-budget figures are *projected* or *modeled* — they rest on validated physics but await fabrication. The boundary between demonstrated and projected is this paper's central scientific discipline.

2. ARCHITECTURE

2.1 Eigenmode Encoding

A glass rod of length L supports longitudinal eigenmodes at:

$$f_n = \frac{nv_{\text{bar}}}{2L}, \quad v_{\text{bar}} = \sqrt{\frac{E}{\rho}}$$

The maximum number of resolvable, thermally stable modes is:

$$n_{\max} = \left\lfloor \frac{1}{2\alpha\,\Delta T + 1/Q} \right\rfloor$$

This depends only on material properties (α, Q) and temperature stability (ΔT), not on rod length. For borosilicate ($\alpha = 3.3 \times 10^{-6}$/K, $Q = 10{,}000$, ± 1 K): $n_{\max} = 9{,}380$.

Each mode carries information up to the Shannon limit:

$$b = \frac{1}{2}\log_2(1 + \text{SNR})$$

At the thermodynamic SNR ceiling of 98.5 dB: $b = 16.4$ bits/mode (derived). At measured bench SNR of 55 dB: $b = 9.2$ bits/mode.

2.2 Perturbation Encoding (Write)

A mass perturbation at position x_0 shifts mode n by:

$$\frac{\Delta\omega_n}{\omega_n} = -\frac{1}{2}\frac{\Delta m \cdot u_n^2(x_0)}{m_{\text{eff}}}$$

Each mode has a different spatial shape $u_n(x)$, so the same mass creates a different shift for each mode. Different mass patterns → different spectral fingerprints → different stored data.

2.3 Interference Recall (Read / Compute)

Drive an array of M resonators with a query spectrum $\{Q_1, \ldots, Q_N\}$. Each element responds with amplitude proportional to the inner product of its stored fingerprint and the query:

$$R_j = \sum_{n=1}^{N} A_n^{(j)} Q_n$$

The maximum responder is the best match — a content-addressable memory implemented in wave physics. This is mathematically equivalent to a Hopfield network [9], with capacity $P_{\max} \approx 0.138\,N$ [10].

PART II

Substrate and Prototype

3. MACRO-SCALE PROTOTYPE

3.1 Hardware Platform

The prototype is a fused silica plate (100×100×1 mm) with four PZT transducers at diagonal corners. Total core materials cost: $38.

Component	Specification
Substrate	Fused silica plate, 100×100×1 mm
Transducers	4× PZT 20 mm disc (SW=TX, NW/NE=RX, SE=spare)
Drive	Raspberry Pi Pico H NCO (3-ch, GP2/GP3/GP4, 126 MHz PIO)
Readout	PicoScope 2204A (2-ch simultaneous, f_s = 781.25 kHz, N = 3968)
Preamp	OPA2134PA (×11 gain, ±9 V supply) on Ch A
Multiplexing	Arduino relay mux (8 channels, 9600 baud)

Signal chain: Pico NCO → 220Ω → TX PZT → plate (acoustic) → RX PZT → preamp → PicoScope.

3.2 Q-Factor Measurement

Method: Ringdown at 35,840 Hz. Drive to steady state, remove excitation, fit exponential decay.

Result (Measured): $\tau = 24.5$ ms, giving:

$$Q_{\text{intrinsic}} = \pi f \tau = \pi \times 35{,}840 \times 0.0245 = 2{,}759$$

Loaded Q: When the TX PZT remains connected (low-impedance path to ground via drive electronics), the plate is heavily damped: $Q_{\text{loaded}} = 152\text{–}241$, $\tau_{\text{loaded}} = 1\text{–}4$ ms. This PZT back-loading is the dominant limitation for temporal applications (Section 6).

3.3 Mode Census

Swept 30–120 kHz with the Pico NCO. **27 modes** identified above 3σ noise floor, all with SNR 42–55 dB. Best mode: 97 kHz at 10,812× noise floor (55.6 dB). Mode spacing is irregular (plate geometry, not 1D rod), consistent with Lamb wave dispersion in a square plate.

3.4 Signal-Path Decomposition

A critical confound in the breadboard topology: the TX→RX path includes both acoustic propagation through the plate and electrical feedthrough via shared ground planes.

Three convergent measurements establish the acoustic fraction:

Method	Acoustic fraction
Re-excitation interference contrast (T2.1)	13.2% (3.4σ)
Ring-up vs. steady-state ratio (T1.2)	~12%
PZT-lifted null test (T4.1)	0% feedthrough confirmed

The signal is ~12% acoustic, ~88% electrical feedthrough. However, the PZT-lifted null test (physically removing the TX PZT from the plate) confirms that **all spectral structure originates acoustically** — feedthrough is flat (no mode peaks). Classification results that depend on spectral *shape* (mode positions, amplitude ratios) are therefore valid: the acoustic component carries the information; the electrical component adds only a DC pedestal.

Null-control battery (E36): Correct enrollment = 4/4 (margin +5.31). Shuffled enrollment = 0/4. Random enrollment = 22% (chance). Separation metric = +12.78. This definitively proves that classification exploits rod/plate-specific spectral structure, not feedthrough artifacts.

3.5 Intermodulation and Cross-Mode Independence

Intermodulation products (T2.2): Drove two modes simultaneously (54,920 + 97,011 Hz). Searched for IM products at $f_1 \pm f_2$, $2f_1 - f_2$, etc. **Result: None detected above noise floor.** The plate is a linear system.

Cross-mode coupling (T2.3): Drove mode A, measured response at mode B frequency. **Result: < 1.1σ.** Modes are independent channels.

Implication: The plate functions as a perfect linear spectral filter bank. This is a *strength* for encoding (orthogonality guarantees zero crosstalk) but means the plate cannot function as a nonlinear reservoir or coherent Ising machine at the bench.

4. SPECTRAL ENCODING AND COMPUTATION

4.1 Pattern Discrimination

Protocol: Drive 4 modes (35,840 / 54,920 / 57,037 / 97,011 Hz) in 4 binary patterns (00, 01, 10, 11). 12 repetitions per pattern, leave-one-out cross-validation.

Result (Measured): 100% accuracy (80/80 trials). On/off contrast: 50–70× per mode. Mean separation: 193σ. Classification is trivial — the patterns occupy non-overlapping regions of feature space.

4.2 Multi-Level Amplitude Encoding

Protocol: Drive each of 4 modes at 8 amplitude levels (50–500 mVpp AWG), 20 repetitions per level. Classify full 4-mode × 8-level patterns using nearest-centroid (Mahalanobis).

Result (Measured):

Mode	Frequency	Min separation	Mean separation	Accuracy
0	35,840 Hz	9.0σ	21.4σ	100%
1	54,920 Hz	23.7σ	80.8σ	100%
2	57,037 Hz	9.0σ	42.0σ	100%
3	97,011 Hz	17.1σ	56.3σ	100%

Combined capacity: $8^4 = 4{,}096$ discriminable patterns = **12 bits** per observation at zero error. Conservative floor (3σ criterion applied to all gaps): 8 bits (256 patterns). This is the first empirical confirmation of L^M capacity scaling, enabled by the measured mode orthogonality (Section 3.5).

4.3 Boolean Computation

Protocol: Encode two modes as binary inputs (on/off). Decode AND, OR, XOR from measured amplitude patterns using pre-scan filtering (strong-mode selection) and per-rod self-response detection.

Result (Measured): 100% fidelity on AND, OR, and XOR across all tested configurations:

Configuration	AND	OR	XOR
2-input (4 patterns)	100%	100%	100%
3-input (AND3/OR3/MAJ/XOR3)	100%	100%	100%
Chained: (A∧B)⊕C	—	—	100% (5/5)

The chained Boolean result demonstrates gate composability without signal regeneration — the output of one gate (a set of frequencies) feeds directly as input to the next. Frequencies are stable discrete identifiers; no analog degradation accumulates between stages.

Noise robustness: All operations remain at 100% accuracy down to 10% drive amplitude (0.2 Vpp, 20× attenuation). Discrimination margin *increases* at lower drive (from +2.5 to +4.2), because template scoring operates on normalized amplitude ratios, not absolute values.

4.4 Content-Addressable Memory

Protocol: Enroll 27 modes as addresses. Query with known patterns. Score by template matching (cross-relay normalization, enrollment-based boost/penalize).

Result (Measured): 100% exact retrieval at all stored patterns. Associative recall via nearest-neighbor in spectral space. Margin: +5.28 (mean), reproducible across 3 independent runs. Temporal stability: 100% over 48 hours, 7 sessions (Wilson 95% CI: [75.7%, 100%]).

PART III

Finite Element Validation

5. CLASSICAL NON-SEPARABILITY OF ACOUSTIC EIGENMODES

5.1 Background

Classical entanglement — non-separability of internal degrees of freedom within a single classical field — was established theoretically by Spreeuw (1998) [13] and measured in optical beams by Kagalwala et al. (2013) [14], following the formalism of Qian & Eberly (2011) [15]. These works demonstrated that the CHSH inequality, conventionally a witness of quantum nonlocality between spatially separated particles, also serves as a non-separability witness for the internal structure of a single classical wave.

We report the first extension of this framework to acoustic eigenmodes in a solid-state resonator.

Explicit disclaimer: No quantum nonlocality is claimed. Bell's theorem is not violated in the EPR sense. The system is a single classical plate; both degrees of freedom (frequency and spatial position) are measured simultaneously at the same location. The CHSH parameter S is used strictly as a mathematical witness of non-separability, following established convention in the classical entanglement literature [13–17].

5.2 Formalism

Consider a plate driven simultaneously at two eigenfrequencies f_1 and f_2, with response measured at two spatial receivers (Ch A = NW, Ch B = NE). The state is described by a 2×2 matrix:

$$M_{ij} = \text{magnitude at frequency } f_i \text{ and receiver } j$$

After row-normalization and Frobenius normalization, this matrix lives in the same Hilbert-space structure as a two-qubit density matrix. Its non-separability is quantified by the concurrence:

$$C = \frac{2\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2}$$

where σ_1, σ_2 are the singular values of M. A separable (product) state has $C = 0$; a maximally non-separable state has $C = 1$.

The CHSH parameter is computed from intensity correlators:

$$S = |E(a_1, b_1) - E(a_1, b_2) + E(a_2, b_1) + E(a_2, b_2)|$$

where $E(a,b) = \frac{I(a,b)+I(a^\perp,b^\perp)-I(a,b^\perp)-I(a^\perp,b)}{I(a,b)+I(a^\perp,b^\perp)+I(a,b^\perp)+I(a^\perp,b)}$

and $I(a,b) = (\hat{a}^T M \hat{b})^2$ is the projected intensity along abstract analyzer angles $\hat{a}$, $\hat{b}$. The four optimal angles are found by numerical optimization (differential evolution). For any separable state, $S \leq 2$; for a maximally non-separable state, $S = 2\sqrt{2} \approx 2.828$ (the Tsirelson bound).

5.3 Multi-Pair Results (E1)

Protocol: 200 trials × 20 averages per trial. Dual-channel simultaneous capture (zero phase jitter). Five mode pairs selected by frequency sweep to maximize spatial contrast (ratio B/A ≪ 1 for one mode, ≫ 1 for the other).

Mode pair (Hz)	S	95% CI	Concurrence	σ above 2.0
34,000 + 70,000	**2.8271**	[2.8271376, 2.8271527]	0.9991	218,744
34,000 + 87,000	**2.8251**	[2.8250816, 2.8251047]	0.9976	137,707
70,000 + 112,000	**2.8239**	[2.8239070, 2.8239366]	0.9968	107,527
34,000 + 80,000	**2.8223**	[2.8222726, 2.8223136]	0.9957	78,594
34,000 + 71,000	**2.8195**	[2.8195104, 2.8195819]	0.9937	45,830

All five pairs violate $S = 2$ at extreme significance. The best pair reaches 99.95% of the Tsirelson bound. All 10/10 block sub-analyses pass independently. The result is not cherry-picked: mode pairs were selected by maximizing spatial contrast (log-ratio between receivers), and all tested pairs with high contrast produce $S > 2.82$.

State matrix (optimal pair, row/Frobenius normalized):

$$M = \begin{pmatrix} 0.707 & 0.014 \\ 0.013 & 0.707 \end{pmatrix}$$

This anti-diagonal structure indicates maximal non-separability: each frequency mode couples predominantly to a different spatial receiver.

5.4 Magnitude-Only Protocol Validation (E2)

Problem: Phase between spatially separated receivers is unstable (42° std for f_1, 19° std for f_2) because one mode has a spatial node at one receiver (signal ≈ 124 counts vs. 63,000 at the other).

Result: Complex tomography gives $C = 0.924$ with CI [0.20, 0.999] (huge variance). Magnitude-only gives $C = 0.999 \pm 0.0000$ (rock-solid).

Conclusion: The magnitude-based formalism of Qian & Eberly (2011) is not merely a convenience — it is the physically correct protocol for spatially separated receivers. Phase is ill-defined at nodal positions. This experiment validates our measurement choice rather than discovering a limitation.

5.5 Temporal Stability (E3)

Seven epochs over 3.5 hours. Temperature range: 22.4–27.0°C.

Metric	Value
Mean S	2.8261
Std S	±0.0003
Max drift from baseline	0.65%
Concurrence (all epochs)	> 0.998

The non-separable state is a stable physical fingerprint determined by plate geometry — not a transient or environmentally sensitive phenomenon. This stability underlies the physical unclonable function (PUF) application.

5.6 Endurance (E11)

16,480,345 continuous drive cycles at 54,920 Hz. Maximum drift from baseline: **0.22%**. No fatigue degradation detected. The acoustic eigenmode spectrum is a non-volatile property of the plate's geometry.

5.7 Three-Mode Extension (E5)

Three simultaneous drive frequencies (35,840 + 54,920 + 97,011 Hz), dual-channel capture.

Metric	Value
Schmidt number K	1.004
σ above $K = 1$	1,184
Concurrence	0.09
Angular spread	6.5°

Interpretation: $K > 1$ is statistically confirmed (>1000σ), proving that the non-separability framework extends to higher dimensions. However, concurrence is low because the NW and NE receivers subtend only 6.5° of angular diversity in mode-shape space. Full 3-mode non-separability requires more spatially diverse receivers (e.g., relay-muxed array or differently positioned PZTs). This is a geometry limitation, not a physics one.

5.8 Computational Utility (E8)

Protocol: Drive the CHSH-optimal mode pair (34 kHz + 70 kHz) in 4 binary patterns (00, 01, 10, 11). Decode at both receivers independently.

Result: 100% accuracy at both Ch A and Ch B. However, the intensity ratios differ:

- Ch A: f_1/f_2 ratio = 0.201 (pattern 11)
- Ch B: f_1/f_2 ratio = 0.361 (pattern 11)
- Ratio of ratios: 0.557 ≠ 1.0 (proves non-separable view)

Interpretation: A separable system ($C = 0$) would show identical ratios at both receivers. The non-separable system provides each receiver with a distinct "view" of the same transmitted state — a physical basis for spatial multiplexing. For analog (multi-level) encoding, this means each receiver resolves different amplitude patterns, extracting more total information than any single channel could.

PART IV

MEMS Design and Scaling

6. RESERVOIR COMPUTING AND TEMPORAL MEMORY

6.1 What the Bench Shows

Three temporal memory architectures were tested. All fail.

Approach	NRMSE / Metric	Root Cause
NARMA-10 (fundamental ringdown)	1.08	τ_{loaded} = 1–4 ms; 140 ms step interval → zero memory
NARMA-10 (intermodulation ring-up)	5.36	IM products genuine but zero persistence
NARMA-10 (active drive modulation)	> 1.0	Active drive masks input history

Physics diagnosis: The bench operates at step_interval/$\tau \approx 100$. Inter-step memory: $\exp(-140/1.4) = 3.7 \times 10^{-44}$ (effectively zero). The plate at bench rates is a memoryless static function, not a reservoir.

6.2 What the Bench Validates

Despite the temporal failure, the bench confirms two reservoir prerequisites:

Lorentzian kernel (Measured): Fine frequency sweep (±500 Hz around 35,840 Hz) reveals a well-resolved resonance: $Q_{\text{loaded}} = 473$, 4.9× dynamic range. The plate implements a Lorentzian nonlinear mapping from input frequency to output amplitude — a valid reservoir kernel.

Spatial diversity (Measured): The H-matrix (26 modes × 2 receivers) has condition number 10.96 and full rank (2). Mode amplitude ratios between receivers span 1.3–11.2×, providing the input-dependent mixing required for echo state computation.

6.3 Simulation with Measured Parameters

Using the measured H-matrix and Q-derived mode dynamics, we simulate reservoir operation at the correct rate regime:

Configuration	NRMSE	Features	Comparison
Round-robin, $G = 3$ + quadratic readout	0.393	384	PASS
Delay-line $D = 15$ + quadratic	0.373	216	PASS
Full embedding $D = 15$	0.351	1,010	PASS
Wiener filter (no plate, reference)	0.241	135	Lower bound
Random ESN (27 nodes, reference)	0.442	—	Upper bound

The plate-parameter reservoir (NRMSE = 0.39) outperforms a random echo state network of equivalent size (0.44), confirming that the measured spatial diversity and Lorentzian kernel provide computational structure. The gap to the Wiener filter (0.24) quantifies the remaining nonlinearity deficit — addressable by MEMS-scale Q enhancement.

6.4 The Scaling Argument

Parameter	Bench	MEMS (projected)
Q_{loaded}	241	9,097
Mode decay time τ	1.4 ms	52 ms
Step interval	140 ms	2 µs
step/τ ratio	100	0.00004
Inter-step memory	10^{-44}	0.99996

The bench proves the **kernel**; MEMS provides the **memory**. The 2.5-million-fold improvement in step/τ converts a memoryless classifier into a functioning temporal reservoir. This is an engineering gap (serial instrumentation, PZT loading), not a physics gap.

7. SCALING LAWS

7.1 Mode Count Is Size-Independent

The key scaling result of Section 2.1 bears repeating: the number of thermally stable modes $n_{\max}$ depends only on α, Q, and ΔT — not on rod length L. A 1 mm MEMS rod supports the same 9,380 modes as the 150 mm prototype (for borosilicate, $Q =$ 10,000, ± 1 K).

7.2 SNR Scales Linearly with Length

The thermodynamic SNR ceiling is:

$$\text{SNR} = \frac{\frac{1}{2} k_{\text{eff}} A^2}{k_B T}$$

The effective spring constant $k_{\text{eff}} \propto E \cdot d^2 / L$, so at fixed drive amplitude: SNR $\propto L$. A 1 mm rod at 1 nm drive has $\text{SNR}_{\text{ceiling}} = 76.7$ dB (derived), giving 12.7 bits/mode. The measured bench SNR (42–55 dB) is below this ceiling; the gap is accounted for by the coupling budget: 88% electrical feedthrough (~9 dB loss of acoustic signal utility), ADC quantization (~20 dB), breadboard pickup (~10 dB), FFT leakage (~5–8 dB).

7.3 Density Scales as $1/L^2$

Since mode count is length-independent and rod volume $\propto L$ (at fixed aspect ratio), areal density scales as $1/L^2$:

Rod length	Bits/rod	Active density	Packed-array density
150 mm (prototype)	153,832	36 Mbit/cm³	7.2 Mbit/cm³
1 mm (MEMS boro.)	119,126	95.1 Gbit/cm³	17.0 Gbit/cm³
0.5 mm (SiO_2)	1.44M	—	1.4 Tbit/cm³

All density figures are *projected*. They assume the full mode count is physically realizable, which requires $Q >$ 1,000 at MEMS scale.

8. MEMS Q-FACTOR MODEL

The central question for MEMS viability: does the quality factor survive miniaturization?

8.1 Five-Mechanism Loss Budget

We model five independent dissipation mechanisms for a 1 mm × 40 µm fused silica rod, vacuum-packaged, with AlN thin-film transduction:

Mechanism	Q_i	Fraction of loss
Material (intrinsic glass)	10,000	91.0%
Anchor (substrate radiation)	207,000	4.4%
Thermoelastic (TED)	2.78×10^6	0.3%
Gas damping (at 1 mTorr)	1.25×10^6	0.7%
Surface loss (50 nm layer)	272,000	3.6%

$$\frac{1}{Q_{\text{total}}} = \sum_i \frac{1}{Q_i} \implies Q_{\text{total}} = 9{,}097$$

Material loss dominates at 91% of the total budget. This is favorable: it means the MEMS-specific mechanisms (anchor, surface, gas) contribute only 9% of the loss, and the Q is determined primarily by the glass material itself — which is well-characterized and consistent across geometries.

8.2 Validation Against Measured Data

The model predicts $Q = 9{,}097$ for MEMS. Our measured values bracket this:

- Plate intrinsic: $Q = 2{,}759$ (fused silica, loaded by multiple PZTs and air)
- Plate (AWG ringdown, single mode): $Q = 7{,}687$–$33{,}960$ (range across modes)
- Rod prototype (borosilicate, epoxied PZT): $Q = 74$–572 (coupling-loss dominated)

The plate's intrinsic $Q = 2{,}759$ is limited by air damping and PZT loading — both eliminated in vacuum-packaged MEMS. The AWG ringdown values reaching $Q >$ 30,000 at some modes confirm that the material itself supports the predicted Q when extrinsic losses are reduced.

9. MEMS DEVICE SPECIFICATION

9.1 Reference Design

Parameter	Value
Rod material	Fused silica
Rod dimensions	1 mm × 40 µm (25:1 aspect ratio)
Transduction	AlN thin-film piezoelectric (200 nm)
Packaging	Wafer-level vacuum (< 1 mTorr)
Array pitch	80 µm (2× diameter)
Layer spacing	1.1 mm

9.2 Per-Rod Performance (Projected)

Parameter	Value
Mode count	9,380
Bits per mode	12.7 (at 76.7 dB SNR)
Bits per rod	119,126
Read time (impulse)	3.8 µs
Read energy	1.7 pJ/bit
Write energy	15 fJ/bit

9.3 Array Architecture (Projected)

A 10×10 mm die at 80 µm pitch contains 15,625 rods in a single layer. With 1.1 mm layer spacing, a 10-layer stack fits in a standard package:

- Rods per die: 156,250
- Bits per die: 18.6 Gbit (2.3 GB)
- Volume: 0.11 cm^3
- Packed-array density: 17.0 Gbit/cm^3

9.4 Energy Budget (Projected)

Operation	Energy
Write (physics-layer)	15 fJ/bit
Read (ADC-dominated)	1.7 pJ/bit
Search (parallel, per rod)	195 fJ
Total system read	~2 pJ/bit

Comparison: DRAM ~3 pJ/bit (access), Flash ~10 pJ/bit (read). CWM is competitive on energy while providing native associative computation.

PART V

Advanced Techniques

10. FABRICATION PATHWAY

All steps use processes in volume MEMS production today:

1. **Fused silica wafer** (200 mm, 500 µm thick)
2. **Deep reactive ion etch (DRIE)** — define rod arrays (Bosch-type, proven for glass MEMS oscillators)
3. **AlN sputter deposition** — 200 nm thin-film piezoelectric transduction layer
4. **Metal patterning** — Mo bottom electrode, Al top electrode (standard lift-off)
5. **Perturbation patterning** — Au dot array via e-beam lithography or focused ion beam (50 nm thick, positioned at calculated antinodes)
6. **Wafer-level vacuum packaging** — getter-sealed cavity (< 1 mTorr)

CMOS integration: Flip-chip bonding to a readout ASIC (ADC array + FFT accelerator + template matching logic). Standard 3D-IC process.

Risk assessment: The individual process steps are mature. The integration — combining glass DRIE with AlN transduction and perturbation patterning on the same wafer — has not been demonstrated for this specific application. The fabrication pathway is the primary technical risk.

11. TECHNOLOGY COMPARISON

Technology	Density (Gbit/cm^3)	Read energy (pJ/bit)	Latency	Compute-in-memory
DRAM	0.5–1	3	10–20 ns	No
3D NAND	10–30	10	25–100 µs	No
ReRAM crossbar	1–5	0.1–1	10–100 ns	Matrix-vector
PCM	1–3	1–10	50–150 ns	Matrix-vector
CWM (MEMS, projected)	**17**	**1.7**	**3.8 µs**	**Associative + Boolean**

CWM's architectural distinction: the medium simultaneously *is* the storage and *performs* the computation. ReRAM/PCM crossbars compute matrix-vector products but cannot perform content-addressable search in a single operation. CWM's physics-layer search (drive + max-response detection) is inherently O(1) in the number of stored patterns.

What CWM is not: CWM is not faster than DRAM or Flash for random access (3.8 µs vs. 10 ns). Its advantage is in *search* and *pattern matching* workloads where conventional architectures require sequential comparison or dedicated accelerators.

12. DISCUSSION

12.1 Evidence Scorecard

Claim	Status	Evidence
Eigenmode encoding in glass plates	**Confirmed**	27 modes, 100% discrimination, 8-bit capacity
Non-separable frequency×space states	**Confirmed**	$S > 2.82$, 5/5 pairs, 200kσ
Non-separability as computational resource	**Confirmed**	Spatial multiplexing, ratio-of-ratios ≠ 1
Boolean computation (AND/OR/XOR)	**Confirmed**	100%, 3-input, composable
Content-addressable memory	**Confirmed**	27 addresses, 100% retrieval
Physical fingerprint stability (PUF)	**Confirmed**	0.65% drift / 3.5h, 16.5M cycle endurance
Plate linearity (mode orthogonality)	**Confirmed**	Zero IM products, < 1.1σ cross-coupling
Temporal reservoir computing (bench)	**Failed**	Rate/Q mismatch (step/τ = 100)
Attention-layer utility (L3)	**Inconclusive**	H converges but indistinguishable from random (rank-2 limit)
MEMS reservoir computing	**Projected**	Simulation NRMSE = 0.39 with measured H
MEMS Q = 9,097	**Modeled**	Five-mechanism budget; awaits fabrication
Density 17.0 Gbit/cm³	**Projected**	Requires measured MEMS Q > 1,000

12.2 What the Failures Teach

The temporal memory failure is not an architecture failure — it is a rate limitation of the bench instrumentation. The serial communication overhead (115,200 baud NCO commands + PicoScope block acquisition) imposes a 140 ms minimum step interval. With $\tau_{\text{loaded}} = 1.4$ ms, inter-step memory is 10^{-44}. At MEMS rates (2 µs step, 52 ms decay), the same architecture provides memory depth > 10 steps.

The attention-layer result (L3) reveals a fundamental geometric constraint: with only 2 spatial receivers, the transfer matrix H is rank-2. Any learnable layer before H can rotate its inputs to exploit any rank-2 matrix equally well — physical structure becomes algebraically invisible (absorption theorem). This identifies a clear hardware requirement: useful attention computation needs ≥ 8 spatial channels.

12.3 Anticipated Objections

"The CHSH result is just geometry." Correct. The non-separability arises from the plate's Chladni patterns — eigenmode shapes that project differently onto spatially separated receivers. This is real physics (measurable, reproducible, computationally useful), not a measurement artifact. That it is "just geometry" makes it *more* reliable, not less: the geometry is fixed, stable, and manufacturable.

"S = 2.83 implies quantum entanglement." It does not. We follow the established classical entanglement framework [13–17] where the CHSH parameter witnesses non-separability of internal degrees of freedom within a single classical system. No space-like separation, no hidden variables, no quantum nonlocality. The mathematical framework is shared; the physical interpretation is entirely different.

"The plate is 88% electrical feedthrough." The spectral *structure* (mode positions, amplitude ratios) is 100% acoustic, confirmed by PZT-lifted null test and the E36 null-control battery. Feedthrough adds a flat pedestal; it does not create or modify spectral peaks. All classification and non-separability results depend on spectral shape, not absolute amplitude.

"Why not just use a digital matched filter?" A digital matched filter requires: (1) storing all templates in electronic memory, (2) streaming each template through a multiply-accumulate unit, (3) comparing all results. For M templates of length N: $O(MN)$ operations, $O(MN)$ memory. CWM performs the same operation in $O(1)$ time (one acoustic propagation) with the templates stored in the physics. The advantage is architectural: zero data movement.

"Q = 2,759 is too low for the density claims." The density projections assume MEMS Q = 9,097 (modeled). The measured plate Q = 2,759 is limited by air damping and PZT loading, both eliminated in vacuum-packaged MEMS. The model's prediction is consistent with published MEMS Q values for fused silica resonators [11, 12]. A single fabricated device would resolve this.

"Reservoir computing is only simulated." Correct. We clearly label it as *modeled*. The simulation uses measured physical parameters (H-matrix, Q) and makes one assumption: that the step interval can be reduced to 2 µs (MEMS ASIC). This is not speculative — it is standard MEMS readout timing. The bench validates the kernel; the projection validates the memory. Together they predict a working reservoir, but the claim remains projected until fabricated.

12.4 Limitations and Open Questions

1. **No MEMS device exists.** All projected performance awaits fabrication. The primary kill criterion is $Q < 1{,}000$ at MEMS scale.
2. **Write-once architecture.** Perturbation patterns are fixed at manufacture. Rewritability paths exist (Section 12 of v18) but are unvalidated.
3. **Temperature sensitivity.** The ± 1 K assumption for 9,380 modes requires active thermal control at MEMS scale. Relaxing to ± 10 K reduces mode count to ~940.
4. **Read speed.** 3.8 µs is slow compared to DRAM (10 ns). CWM targets search/pattern-matching workloads, not random access.
5. **2-receiver geometry limits.** Current experiments use 2 spatial channels. Full exploitation of the mode-space diversity requires 8+ receivers.
6. **Endurance of perturbation patterns under thermal cycling** is untested beyond room-temperature stability.

12.5 Related Work

Classical entanglement: Spreeuw (1998) [13] proposed the concept; Kagalwala et al. (2013) [14] measured $S > 2$ in optical beams; Qian & Eberly (2011) [15] developed the intensity-based formalism; Töppel et al. (2014) [16] and Aiello et al. (2015) [17] extended the framework. All prior demonstrations are in optics. This work extends classical non-separability to acoustic eigenmodes in a solid resonator.

Photonic reservoir computing: Larger et al. (2012) [18], Brunner et al. (2013) [19] demonstrated optical cavities as reservoir nodes. CWM proposes the same architecture in acoustic MEMS — slower speed but higher Q and lower energy.

MEMS computing: Mahboob et al. (2011) [20] demonstrated coupled MEMS oscillators for logic. Their approach uses nonlinear coupling between oscillators; CWM uses linear superposition within a single multimode resonator plus spatial diversity.

PART VI

Outlook

13. CONCLUSION

We have demonstrated that acoustic eigenmodes in glass plates constitute a viable physical substrate for information encoding and computation. The central results are:

1. **Non-separable frequency×space correlations** quantified by $S = 2.827$ (99.95% of Tsirelson bound) at $>$200,000σ significance — the first measurement of classical entanglement in an acoustic solid-state system.
2. **Complete spectral encoding validation:** 27 modes, 8-bit multi-level capacity, 100% Boolean computation, content-addressable memory — all on a $38 prototype.
3. **Honest identification of the current limit:** temporal reservoir computing fails on the bench due to a 2.5-million-fold rate/Q mismatch, not due to physics. The mea-

sured Lorentzian kernel and spatial H-matrix predict NRMSE = 0.39 at MEMS rates.

4. **A clear fabrication pathway** using processes in volume production today, with a rigorous Q-factor model ($Q = 9{,}097$) predicting that miniaturization preserves the physics.

The non-separability result, while classical, has immediate practical utility: it provides spatial multiplexing (different receivers decode different information from the same drive), enables physical unclonable functions (stable fingerprints, 16.5M cycle endurance), and underlies the mode-diversity that makes spectral encoding work. This is computationally useful entanglement in a $38 glass plate.

What remains is to build the MEMS device and measure it.

REFERENCES

[1] J. Backus, "Can programming be liberated from the von Neumann style? A functional style and its algebra of programs," *Commun. ACM* **21**, 613–641 (1978).

[2] W. A. Wulf and S. A. McKee, "Hitting the memory wall: Implications of the obvious," *ACM SIGARCH Comp. Arch. News* **23**, 20–24 (1995).

[3] F. Cai et al., "A fully integrated reprogrammable memristor–CMOS system for efficient multiply–accumulate operations," *Nat. Electron.* **2**, 290–299 (2019).

[6] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.* **27**, 379–423 (1948).

[7] Lord Rayleigh, *The Theory of Sound*, 2nd ed. (Macmillan, London, 1894), Vol. 1.

[8] G. Sauerbrey, "Verwendung von Schwingquarzen zur Wägung dünner Schichten und zur Mikrowägung," *Z. Phys.* **155**, 206–222 (1959).

[9] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Natl. Acad. Sci. USA* **79**, 2554–2558 (1982).

[10] D. J. Amit, H. Gutfreund, and H. Sompolinsky, "Storing infinite numbers of patterns in a spin-glass model of neural networks," *Phys. Rev. Lett.* **55**, 1530–1533 (1985).

[11] F. Ayazi and K. Najafi, "High aspect-ratio combined poly and single-crystal silicon (HARPSS) MEMS technology," *J. Microelectromech. Syst.* **9**, 288–294 (2000).

[12] J. R. Vig and Y. Kim, "Noise in microelectromechanical system resonators," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.* **46**, 1558–1565 (1999).

[13] R. J. C. Spreeuw, "A classical analogy of entanglement," *Found. Phys.* **28**, 361–374 (1998).

[14] K. H. Kagalwala, G. Di Giuseppe, A. F. Abouraddy, and B. E. A. Saleh, "Bell's measure in classical optical coherence," *Nat. Photonics* **7**, 72–78 (2013).

[15] X.-F. Qian and J. H. Eberly, "Entanglement and classical polarization states," *Opt. Lett.* **36**, 4110–4112 (2011).

[16] F. Töppel, A. Aiello, C. Marquardt, E. Giacobino, and G. Leuchs, "Classical entanglement in polarization metrology," *New J. Phys.* **16**, 073019 (2014).

[17] A. Aiello, F. Töppel, C. Marquardt, E. Giacobino, and G. Leuchs, "Quantum-like nonseparable structures in optical beams," *New J. Phys.* **17**, 043024 (2015).

[18] L. Larger, M. C. Soriano, D. Brunner, L. Appeltant, J. M. Gutiérrez, L. Pesquera, C. R. Mirasso, and I. Fischer, "Photonic information processing beyond Turing: an optoelectronic implementation of reservoir computing," *Opt. Express* **20**, 3241–3249 (2012).

[19] D. Brunner, M. C. Soriano, C. R. Mirasso, and I. Fischer, "Parallel photonic information processing at gigabyte per second data rates using transient states," *Nat. Commun.* **4**, 1364 (2013).

[20] I. Mahboob, E. Flurin, K. Nishiguchi, A. Fujiwara, and H. Yamaguchi, "Interconnect-free parallel logic circuits in a single mechanical resonator," *Nat. Commun.* **2**, 198 (2011).

[21] Festi et al., "Universality of vibrational spectra of stable glasses," *Phys. Rev. X* **16**, 021021 (2026).

[22] W. Schirmacher, G. Ruocco, and T. Scopigno, "Acoustic attenuation in glasses and its relation with the boson peak," *Phys. Rev. Lett.* **98**, 025501 (2007).

[23] M. Tierce, *Coherent Wave Memory: The Full Story* (companion volume), github.com/miketierce/cwm-book (2026).

APPENDIX A: SNR AND DENSITY SCALING DERIVATION

The thermodynamic SNR for a longitudinal mode of a cylindrical rod:

$$\text{SNR} = \frac{\frac{1}{2}k_{\text{eff}}A^2}{k_B T}$$

where $k_{\text{eff}} = \frac{E\pi d^2}{4L} \cdot \frac{n^2\pi^2}{2}$ for mode n. At the fundamental ($n = 1$), for a 150 mm × 6 mm borosilicate rod ($E = 63$ GPa, $\rho = 2{,}230$ kg/m³):

$$k_{\text{eff}} = \frac{63 \times 10^9 \times \pi \times (6 \times 10^{-3})^2}{4 \times 0.15} \times \frac{\pi^2}{2} = 5.86 \times 10^7 \text{ N/m}$$

At 1 nm drive: $E_s = \frac{1}{2} \times 5.86 \times 10^7 \times (10^{-9})^2 = 2.93 \times 10^{-11}$ J.

$$\text{SNR} = \frac{2.93 \times 10^{-11}}{4.14 \times 10^{-21}} = 7.07 \times 10^9 \quad (98.5 \text{ dB})$$

This is a *derived* thermodynamic ceiling, not a measured instrument reading. Measured bench SNR is 42–55 dB; the 43–56 dB gap is accounted for by the coupling budget (Section 3.4).

Density scaling. For a packed array at pitch $p = 2d$ and layer spacing $h = L + 2.5d$:

$$\rho_{\text{packed}} = \frac{n_{\max} \times b}{p^2 \times h}$$

At 1 mm, 40 µm diameter, 9,380 modes, 12.7 bits/mode: $\rho_{\text{packed}} = 119{,}126/(80 \times 10^{-6})^2 \times (1.1 \times 10^{-3}) = 17.0$ Gbit/cm³ (projected).

APPENDIX B: Q-FACTOR MODEL DETAILS

Material Loss

Borosilicate: $Q_{\text{mat}} = 10{,}000$ (literature, confirmed by macro ringdown). Fused silica: $Q_{\text{mat}} = 100{,}000$ [11]. CWM reference design uses fused silica at $Q_{\text{mat}} = 10{,}000$ (conservative).

Anchor Loss

For a beam resonator with clamped-free boundary conditions [11]:

$$Q_{\text{anchor}} = C\left(\frac{L}{d}\right)^3$$

At 25:1 aspect ratio: $Q_{\text{anchor}} \approx 207{,}000$. Contributes 4.4% of loss.

Thermoelastic Damping (TED)

$$Q_{\text{TED}}^{-1} = \frac{E\alpha^2 T}{\rho C_p} \cdot \frac{\omega\tau_{\text{th}}}{1 + (\omega\tau_{\text{th}})^2}$$

where $\tau_{\text{th}} = d^2/(\pi^2\kappa)$ is the thermal relaxation time. For 40 µm fused silica at 2.66 MHz (fundamental of 1 mm rod): $Q_{\text{TED}} = 2.78 \times 10^6$.

Gas Damping

At pressure P in a vacuum cavity:

$$Q_{\text{gas}} = \frac{\rho d\omega}{4P}\sqrt{\frac{\pi k_B T}{2m_{\text{gas}}}}$$

At 1 mTorr (N_2): $Q_{\text{gas}} = 1.25 \times 10^6$.

Surface Loss

$$Q_{\text{surface}}^{-1} = \frac{4\delta}{d} \cdot Q_{\text{bulk}}^{-1}$$

With surface damage layer $\delta = 50$ nm and bulk dissipation $Q_{\text{bulk}} = 100{,}000$ for fused silica: $Q_{\text{surface}} = 272{,}000$ at $d = 40$ µm.

Combined

$$Q_{\text{total}}^{-1} = 10^{-4} + 4.8 \times 10^{-6} + 3.6 \times 10^{-7} + 8.0 \times 10^{-7} + 3.7 \times 10^{-6} = 1.10 \times 10^{-4}$$

$$Q_{\text{total}} = 9{,}097$$

APPENDIX C: EXPERIMENT GUIDE (REPRODUCIBILITY)

C.1 Reproducing the Non-Separability Measurement

Materials (~$50):

- Any glass plate or rod ($Q > 500$)
- 3× PZT disc (20 mm): 1 TX, 2 RX at different positions
- Dual-channel oscilloscope (e.g., PicoScope 2204A)
- Signal generator (square wave, 30–120 kHz)
- 2× 220Ω resistor
- ~20 lines of Python

Protocol:

1. Attach TX PZT to one corner, two RX PZTs at non-adjacent corners
2. Connect RX PZTs to two scope channels (simultaneous capture)
3. Frequency sweep: find two modes where ratio B/A is very different (one mode strong on Ch A, another strong on Ch B)
4. Drive both modes simultaneously for 3 s (steady-state)
5. Capture FFT at both channels (20 averages, 200 trials)
6. Extract peak magnitude at each drive frequency in each channel → build 2×2 state matrix
7. Row-normalize, Frobenius-normalize
8. SVD → concurrence C
9. Optimize CHSH angles via numerical optimization
10. Bootstrap for confidence interval

Expected result: For any plate with eigenmodes having different spatial distributions at the two receiver positions, $S > 2$ is expected. Higher spatial contrast (log-ratio) → higher S.

Data and code: All scripts, raw data, and analysis code are available at github.com/miketierce/cwm.

C.2 Reproducing Boolean Computation

Additional materials: Arduino Nano (relay mux), 8-channel relay module.

Protocol:

1. Enroll each mode's self-response at all frequencies (pre-scan)
2. Classify modes as strong/weak by geometric-mean threshold
3. Drive pairs of modes in ON/OFF combinations
4. Detect: rod responds if self-response > 50% of its weakest strong peak
5. Decode: AND = both respond, OR = either, XOR = exactly one

Expected result: 100% accuracy for mode pairs with spectral overlap ≤ 40%.

C.3 Reproducing Multi-Level Encoding

Protocol:

1. Drive each mode at L amplitude levels (equally spaced, 50–500 mVpp)
2. 20 repetitions per level per mode
3. Classify using nearest-centroid (per-mode Mahalanobis distance)

Expected result: 100% accuracy for $L = 8$ (per mode) if minimum inter-level separation $> 3\sigma$.
