

Memory by Subtraction

A class-encoded memory architecture, its algebra, its read head, and the experiments that would kill it

Matthew “Bear” Schulz · Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com *Preprint, not peer reviewed. Speculative synthesis, fenced throughout. CC BY 4.0. Draft v1, 2026-08-07.*

In three words: Holonomy-Formatted Memory.

A record whose bit is a *class* rather than a magnitude — a count of holes cut into a solid, deposited on a loop, and read back as an integer. If you only read one line of this paper, read that one; everything else is what it costs, what it buys, and what would prove it wrong.

BOTTOM LINE UP FRONT

Every bit in every deployed memory is a magnitude. A retardance, a resistance, a charge, a reflectivity. Magnitudes decay, and every archival lifetime figure in existence is an extrapolation of how slowly one particular magnitude slides below a readable threshold.

This paper asks what the other column costs and buys. **Store an integer instead — a count of holes.** A count does not slide. It is or it isn't.

Four things follow, and three of them are somebody else's:

1. **The design gate.** *Name your structure group before you fabricate.* Abelian and discrete → exactly computable, no convention to argue about. Non-abelian, or reduced modulo a symmetry → defined only up to a choice you must publish. This sorts memory proposals on paper, before a fab budget.
2. **The algebra of the cut, which is the contribution.** A sealed cavity adds a 2-sphere and gives you an abelian $\mathbb{Z}^n$ readout — an **address**. A through-tunnel adds a handle and gives you a free non-abelian readout — a **sentence**. The map between them is the abelianisation. **You choose your readout algebra at write time**, which no storage scheme currently does.
3. **The write head already exists.** Nanovoids in fused silica, characterised twenty years ago, with period, count and connectivity already on dials.
4. **The read head is a sixty-year-old theorem.** For a planar domain with h holes the heat-trace constant term is $a_0 = (1 - h)/6$. You can hear the number of holes. You cannot hear the shape — and that turns out to be a security property rather than a limitation.

What we are not claiming. Not the fabrication. Not the read mathematics. Not density — we lose on density and say so early. Not a comparison against anyone's product. And nothing here is measured by us: this is theory and synthesis with a specified experimental programme, and Section 9 lists the five experiments that would end it.

Keywords: topological memory · class-encoded storage · Betti numbers · Euler characteristic · fused silica · femtosecond laser writing · nanovoid · microexplosion · homotopy · fundamental group · abelianisation · structure group · gauge · heat trace · Kac · isospectral domains · persistent homology · physically unclonable function · chirality · vortex beam · acoustic vortex · Chladni figures · archival storage · fidelity map

SECTION 0 — WHERE THIS CAME FROM

(Author's voice. Provenance, not proof. Skip it if you want the architecture.)

The design did not begin as a device. The crystal began as an attempt at artificial life in its own right, and it became more interesting as a vessel.

When I built the crystal Matriarch to run the ALife wall, she used sliced, Chladni-style chiral patterns and merged occlusion-based memory with light projected through the slices, integrating them into a combined output state. The little defects were watermarks and shared ancestry. She pilots **passive pucks**, and the word matters: a mote carries distributed intelligence, and our in-body pucks carry none. The parent holds the intelligence, the power, the timing and the map; the puck holds a resonant identity and does as the field tells it. That is where the idea of *driving* through a substrate on a loop walk started earning its keep, because that is what we do everywhere else: quaternions and holonomy, a path and what comes home changed.

And the daughters have a second job that is really a memory job. They are an **in-situ, radiation-hardened recording system**. Deliver whatever they were sent to deliver — then *stay*, and keep recording. A dosimeter that reports while the dose is being delivered rather than after the sample is recovered, and that keeps its record through the radiation it is there to measure. The value is not the delivery; the value is that **the record is written where the event happens, at the moment it happens, and survives the environment that produced it**. That is a memory requirement before it is a medical one, and it is the requirement this whole paper is about. (*The engineering for that sits in a companion paper and its numbers stay there.*)

Why a class-encoded bit is the right bit for that job is the same argument as Section 1, one level down. A charge-based bit is a pile of charge, and an ionising strike adds or removes charge — the failure is direct. A winding, or a hole, is not a pile of anything. **Local damage cannot reach a non-local invariant**, so the generic strike deposits heat the lattice shrugs off and leaves the count untouched. A recorder that has to survive the thing it is recording wants a bit that fails by discrete events rather than by accumulation.

And then the negative result did the work. A companion paper reported that a crystalline substrate which *can* inherit — a topological register written by merging and mutated by fracture — does not, turned loose with nobody steering, climb the open-ended-evolution wall. That wall did not build the Matriarch. **It told us what she was for.** Every property that disqualifies the crystal as artificial life — no autonomy, faithful copying, silence until read, an inability to self-direct — is exactly the property list for a device somebody else pilots.

Read that list again with a different job in mind and it does not change. No autonomy. Faithful copying. Silence until read. Cannot self-direct. Restores to a setpoint rather than wandering off one. That is not a description of a disappointing organism. **That is the specification for a memory.** The same failure told us twice what the crystal was for, and this paper is the second answer.

So: what would that look like?

Holonomy likes quiet substrates and quantised integers. In 2012 I drew some towers that showed up again fourteen years later as the Log and Anti-Log Towers of the Integer Forest — and that example is a holonomy loop driven through pure mathematical structure. Quiet. No curl. No noise. The number is calibrated at the true zero of the competence surface, because the mathematics has no measurement error in it.

If I can drive a pure number to a geometric address, that felt like memory storage. And by then we had built a semantic language on top of it — the resonant codex and the eigenstate dictionaries — and I had three kinds of loop: a LONG loop, a SHORT loop, and a Subway Interchange.

So this paper is the output of one question: what is possible if you treat the journey as a loop that *deposits a winding*?

Two things have to be true at once, and that symmetry is the whole architecture. **I need a winding that closes a loop. And I need the loop it is stored in to close its own loop.** That, plus a substrate worth writing into, is a trifecta.

I did not start with an agenda beyond exploring the idea. I picked glass and gelatin as real, testable options, and honestly the early modelling shows pros and cons to both — which is why Figure 5 is a matrix rather than a recommendation.

I took holonomy to the chipmakers’ metallurgy sandbox and found out it is not my scene. That was a fit problem rather than a physics problem, and the companion paper says so plainly. **This is where I get to invite memory storage into my sandbox, with my holonomy rules.**

1. TWO WAYS TO STORE A BIT

The distinction that organises this paper is physical, and it splits every storage scheme ever built into two families.

	magnitude-encoded	class-encoded
the stored quantity	a continuous parameter — retardance, resistance, charge, reflectivity	an integer — a Betti number, a winding number, a component count
how it fails	slides below a readable threshold	a discrete topological event, or nothing
what damage looks like	indistinguishable from drift or reader miscalibration	a quantised, identifiable change
lifetime claim rests on	extrapolated decay of the parameter	whether a discrete event has happened
reader agreement	needs a shared threshold calibration	integer count, no threshold
density	excellent	poor

Essentially all deployed storage lives in the left column, and that is not a flaw. It is what makes those media dense and fast, and density is the property the world actually needs. The finest archival medium anyone has built — femtosecond-written birefringent nanogratings in fused silica, with extrapolated room-temperature lifetimes in the billions of years — lives there too, and it is superb work. We cite it with respect and we are going a different direction.

This paper is an attempt at the right column. It will lose on density and it should say so on page one. What it offers instead is a readout with no threshold to calibrate and a failure mode that announces itself in the same units as the data.

Worth noticing in passing: the slow-axis orientation of a nanograting is an angle mod π — already a class-like quantity. The two columns are not enemies, and the most interesting media will probably use both.

And the choice of column is not the only choice. Figure 5 lays out six candidate substrates against the four things that actually decide an architecture: the structure group (is the readout well-posed), the write method (what can you cut), the wave operator (which Betti numbers are audible), and the floor and ceiling. It is deliberately a *matrix* rather than a recommendation.

Figure 5 · The substrate matrix — pick the algebra, then pick the material

substrate	structure group	write	wave operator (what you can hear)	floor	ceiling	the reason to pick it
fused silica (rigid, amorphous)	NONE (amorphous)	fs microexplosion 30-100 nJ	elastic → vector modes: b_1 , maybe b_2	void must exceed imaging resolution	viscous flow closes the void	the archive, no point group, so nothing to gauge-fix
flexible polymer or hydrogel	NONE	focused ultrasound cavitation	membrane → SCALAR Laplacian: Kac exact	cavitation threshold and gel set time	creep: strain shifts frequency	the right operator, and it absorbs kinetic damage and returns to shape
single crystal	the point group (12 / 24 / ...)	fs writing, doping, defects	anisotropic — direction carries information	the reduction margin (7,400* for Mg twin)	the gauge problem (951 → 2207)	richest alphabet: the instrument, not the archive
porous / photo-sensitive glass	NONE	expose + etch + sinter 1050-1150 °C	as fused silica	etch selectivity	sintering closes fine structure	connectivity control is already routine here
gas- or vacuum-filled cavity	NONE	the boundary is the payload	acoustic → SCALAR only (b_1 ; no vector modes)	pressure sets Q and the SNR	leak, and the topology is gone	cheapest ring, weakest read; pressure is a sensitivity, not a feature
memristor crossbar	H_1 of the graph $\mathbb{Z}^n$ — abelian	electrical, fast, rewritable	no wave read — electrical cycle sum	per-cell drift	endurance	high changeover; the sidecar, not the archive

Read it left to right: the structure group decides whether the readout is well-posed, the write method decides what you can cut, and the WAVE OPERATOR decides which Betti numbers you can hear. A membrane obeys the scalar Laplacian, so Kac's theorem applies exactly. A stiff plate obeys the biharmonic operator and the theorem must be re-derived. An elastic solid carries vector modes, which is the sector b_2 lives in. Rigidity is not a comfort choice — it selects the mathematics.

Figure 5. The substrate matrix. Read left to right: the structure group decides whether the readout is well-posed, the write method decides what you can cut, and the wave operator decides which Betti numbers you can hear. Rigidity is not a comfort choice — it selects the mathematics.

2. THE GATE: NAME YOUR STRUCTURE GROUP

Before any of the architecture, the rule that generated it.

Name your structure group before you fabricate anything. If it is abelian and discrete, your readout is exactly computable and there is no convention to argue about. If it is non-abelian, or you reduced modulo a symmetry, your readout is defined only up to a choice — and you must publish the choice, or the number is not a number.

This is not philosophy. It was paid for in the companion paper, on real crystallographic data. Composing symmetry-reduced rotations around a loop moved 201 map cells into different defect classes between two computers running identical code on identical bytes, and moved a boundary-seam count from 951 to 2207 across five equally defensible tie-break rules. The same construction on $\mathbb{Z}/2$ is path-independent to machine precision across every route tested: on the integers under multiplication, the Liouville sign is the parity of the subway distance on the exponent lattice — how many exponent changes it takes to get from one integer to another. (The same quantity appears elsewhere as the ℓ_1 or taxicab metric; we use the subway because that is the picture, and we note the equivalence rather than adopt the name.)

And this is the technical argument for glass, which is otherwise chosen for being cheap and everywhere.

An amorphous solid has no point group. No lattice, no quotient, no representative to choose. Fused silica passes the structure-group gate *trivially, because it has no structure group*. Anything topological found in it was put there, so signal-to-background is perfect because the background is empty. The usual ordering inverts: **the disordered medium is the well-posed one.**

The amazing alternative, named as the design rule requires: a single crystal buys you an anisotropic medium where direction itself carries information — six faces, six blind spots, per-face read rules the material enforces rather than the encoder. That is a richer alphabet than glass can offer. It also costs you the entire gauge problem. **Glass for the archive, crystal for the instrument.**

3. THE CUT CHOOSES THE ALGEBRA

This is the contribution, and it fits in a table.

Take a solid block. Remove material. What you can store depends entirely on **how** you remove it.

what you cut	complement gains	invariant lives in	structure group	role
a sealed cavity	a 2-sphere	π_2	$\mathbb{Z}^n$ — abelian, discrete	address
a through-tunnel	a handle	π_1	free group on g — non-abelian	sentence
two linked tunnels	handles with linking	π_1 + linking number	non-abelian, abelian invariant on top	checksum
a knotted tunnel	a knotted handle	the knot group	non-abelian, large	alphabet

Figure 1 · The cut chooses the algebra




Figure 1. The cut chooses the algebra. Four ways to remove material from the same block; each adds a different hole to the complement, and the hole decides which homotopy group holds the data. The map between the first two is the abelianisation.

Read that against the gate and it says something specific:

Cavities give you an address. Tunnels give you a sentence.

Sealed cavities generate a *free abelian* readout: the invariant is a count, order does not matter, no representative to choose, no path-ordering to fix. **Passes the gate.** That is the commutative address space — the exponent lattice, physically realised.

Through-tunnels generate a *free non-abelian* readout: LRRL is not LRLR. **Fails the gate for addressing, and that is exactly the point** — order-dependence is what carries syntax.

And the map between them is the **abelianisation**. Forget the order of a tunnel word and you land back in the cavity lattice. **The same block of glass carries the address and the procedure, and which one you read is a choice of quotient, not a choice of hardware.**

No optical storage scheme currently chooses a readout algebra at write time. They choose a modulation depth.

3.1 Closure is the admission rule

A path that does not close has no holonomy — holonomy is defined only on closed loops. So closure is not a nice property of a well-formed word. **It is the criterion for being a word at all.** An open stroke is a scribble.

That has a cost and a benefit, both countable. Counting closed walks on a lattice:

word length	all 2D words	closed 2D	bits	closed 3D (L, R, i)	bits
4	256	36	5.17	90	6.49
8	65,536	4,900	12.26	44,730	15.45
12	16,777,216	853,776	19.70	32,496,156	24.95

At length 12 in two dimensions, **only 5.1% of possible words are words**; in three dimensions, 1.5%. That rejection ratio is error detection, performed by geometry before any checksum exists — a corrupted word almost certainly fails to close. And the codebook is still large: 853,776 legal words at length 12, against a semantic codex of 432.

(That the closed-word count at length 4 is exactly 36, the size of the chiral winding codec’s alphabet, is a coincidence. We flag it as one and do not build on it.)

3.2 Why the path alphabet needs a third letter

*(An echo worth noting and not building on: neo-Riemannian music theory generates motion on the Tonnetz — which has the topology of a torus — with exactly three operations, **L, P and R**, whose cycles are closed loops with winding numbers (Tymoczko, Science 2006). Three generators, a torus, closed loops, winding. It is a striking coincidence of notation and a genuine echo of structure, and it is **not evidence**: LPR are three involutions acting on triads, a different object from a lattice walk. We record it because someone will notice it, and we would rather name it than be caught by it.)*

L and R are **translations** — infinite order, you can keep going. A pivot is a **rotation** — finite order, it returns. Mixing infinite-order translations with a finite-order rotation is how you build a **space group**, which is the natural home for an address on a lattice. Two practical consequences: L and R alone abelianise to $\mathbb{Z}^2$, a planar address, and a block of glass needs $\mathbb{Z}^3$; and the third symbol makes the alphabet ternary, matching the $\{-1, 0, +1\}$ cell drive the existing tool implementations already use.

4. HOW TO CARVE — AND IT IS SOMEBODY ELSE’S WRITE HEAD

Nanovoids in fused silica are established physics, characterised twenty years ago. Single femtosecond pulses of **30-100 nJ at 800 nm, 180 fs**, at power *below* the self-focusing threshold, form a void by shock and rarefaction, leaving a nanovoid surrounded by a shell of shock-compressed material. The shell is measurably denser — **2.201 → 2.645 g/cm³**, about 20% densification, confirmed by the shell’s higher reactivity under wet etching.

Figure 2 · Where in the write parameter space the payload lives




Figure 2. Where in the write parameter space the payload lives. Type I is index change, Type II is the birefringent nanograting regime that current archival storage uses, Type III is voids by microexplosion — the regime this architecture needs and the one usually avoided as damage.

Three facts that matter for this architecture:

1. **Voids are the standard high-energy regime, not an exotic edge case.** The literature sorts the regimes: isotropic index change at moderate energy (Type I), birefringent nanogratings at intermediate energy (Type II — the magnitude regime), and **voids by microexplosion at high energy (Type III)**. We are asking for the regime everyone else avoids as damage.
2. **Void arrays self-organise and the organisation is controllable.** Periodically aligned nanovoids form along the beam axis, with period, count and total chain length set by laser power, pulse number and focal position. And in two-beam vortex interference, the **number and rotation of the filaments are set by the topological charges of the two beams** — the symbol index is already a machine parameter.
3. **Connectivity control exists, in a different substrate.** Hollow microchannels of arbitrary length and configuration are written routinely in glass, including the deliberate joining of channels, and at the small end a single nano-void of **37 nm** with **the number of voids controlled by peak intensity** and **nanofluidic channels formed by connecting single nano-voids into continuous structures**.

4.1 The substrate mismatch, stated because it reorders the programme

The routine connectivity work is done in *photosensitive or porous* glass, using exposure followed by heat treatment, wet etching, and sintering near 1050–1150 °C. That is a different material and a different mechanism from microexplosion voids in pure fused silica.

The retention argument belongs to fused silica. The connectivity capability currently belongs to etched porous glass. Whether controlled connectivity transfers to the archival substrate is the first experiment, ahead of everything else in Section 9.


4.2 The acoustic partner

Optical vortex beams are the mature route. The acoustic one is under-developed and deserves attention, because focused ultrasound generates cavitation in soft hydrogels routinely — **and a cavitation bubble in gel is an excised void**. Gelatin is transparent, castable, costs nothing, and you can watch the write happen with a camera. It is the cheapest possible demonstrator of the whole architecture and it needs no laser and no clean room.

4.3 Three write heads, what each is for, and whether the speed is a wall

“Slow” without a number is useless, so here are numbers. All of them are order-of-magnitude, none is measured by us, and every assumption is on the page so the estimate can be argued with rather than believed.

Assumptions: roughly 10 voids and 20 bits per symbol (the length-12 three-dimensional figure of Section 3.1), one void per pulse, no beam-steering overhead, no failed writes.



THE MOON CHECK

The Apollo Guidance Computer's entire core rope was 72 KB. At the slowest optical figure above that is **three seconds** of writing; with 60 channels, **fifty milliseconds**. A 4096-bit key set is instantaneous. A bacterial genome is 40 seconds, or under a second parallel.

And the wall is real where it should be: 1 TB takes 463 days single-focus, 7.7 days at 60 channels. This is not a medium for a video library. It is a medium for the thing that must still be readable when nobody remembers what wrote it — keys, a bootloader, a genome, a codebook, a constitution.

Figure 6. Three write heads and what each is actually for. Throughput is dominated by parallelism, not by the physics of a single event.

head	voxel	rate	density	needs	the upside
electron beam	~20 nm	~0.2 kB/s	~4 PB/cm ³	vacuum, thin sample	density nothing else approaches
fs laser, 1 focus	~1 μm	~25 kB/s	~31 GB/cm ³	transparent bulk	the practical middle
fs laser, 25 foci	~1 μm	~625 kB/s	—	+ an SLM	<i>demonstrated</i>
fs laser, 60 channels	~1 μm	~1.5 MB/s	—	+ a fibre array	<i>validated</i>
ultrasound, 10 MHz	~150 μm	~0.2 kB/s	~10 kB/cm ³	nothing	writes into opaque, sealed or living media

Parallelism is the whole story on throughput, and it is not hypothetical. Twenty-five parallel beams have been driven from one femtosecond source, giving an *effective* repetition rate of 500 kHz without restrictive scan speeds, and a sixty-channel fibre array has been validated for parallel direct writing with inter-channel error held inside 20 nm. So the single-focus figure of 25 kB/s — about four times a 56k modem — is the *floor*, not the ceiling. **At sixty channels it is 1.5 MB/s**, which is CD-ROM territory.

The electron beam’s upside is density and only density. A 20 nm voxel is four orders of magnitude smaller in volume than an optical one, so the bits-per-cubic-centimetre figure is enormous. It is also serial, slow, and needs a vacuum and a thin sample — a shared-facility instrument, not an archive writer. **If bits-per-volume is your binding constraint, this is the head.**

The acoustic head’s upside is unique and it is not speed. Ultrasound is the only one of the three that needs *no optical access and no vacuum*. It writes into media that are **opaque, already sealed, or alive** — and it costs a few hundred dollars in parts rather than a few hundred thousand. Its voxel is enormous and its density is dreadful, and neither of those is the point. **The acoustic head is how you write into something you cannot see into.** It is also, per Section 4.2, the cheapest possible demonstrator of the whole architecture.

Is the speed a practical wall?

Not for what this medium is for, and the honest way to show that is to write down what actually needs writing.

The Apollo Guidance Computer’s entire core rope was **72 KB**. At the *slowest* optical figure above, that is **three seconds** of writing. At sixty channels, **fifty milliseconds**. A 4096-bit key set is instantaneous. A bacterial genome is forty seconds single-focus, under a second parallel. The King James Bible is three minutes, or three seconds parallel.

And the wall is real exactly where it should be. One terabyte takes 463 days at a single focus and 7.7 days at sixty channels. **This is not a medium for a video library, and it should stop pretending it might be.**

It is a medium for the thing that must still be readable when nobody remembers what wrote it: keys, a bootloader, a genome, a codebook, a constitution. Those are kilobytes and megabytes, not terabytes — and at that scale the write time disappears entirely against a retention argument measured in centuries.

5. THE READ HEAD IS A THEOREM

Excision stores Betti numbers — b_0 components, b_1 tunnels, b_2 cavities. There are two ways to read them and one of them is free.

The direct route is X-ray micro-tomography plus persistent homology. Mature, off the shelf, and it computes exactly the three numbers this architecture stores. No new instrument, no calibration curve, no threshold to argue about: a tomogram either resolves a tunnel or it does not, and if it does, the count is an integer.

The cheap route is a sixty-year-old theorem. For a smoothly bounded planar domain Ω with h holes, the heat trace has the asymptotic expansion

$$Z(t) = a_{-1}/t + a_{-1/2}/\sqrt{t} + a_0 + a_{1/2}\sqrt{t} + \dots$$

and the constant term is

$$a_0(\Omega, h) = (1 - h)/6$$

Each hole contributes exactly $-1/6$. Kac conjectured this in 1966 by approximating smooth domains with polygons; **McKean and Singer proved it correct one year later**. The area is audible, the perimeter is audible, and **the number of holes is audible**.

Figure 3 · The read head is a theorem — and its limit is a feature




Figure 3. The read head. Left: each hole contributes exactly $-1/6$ to the heat-trace constant term (Kac 1966; McKean-Singer 1967). Right: the spectrum does not determine the shape (Gordon-Webb-Wolpert 1992) — which makes the payload readable without the geometry being recoverable.

Ring the substrate, take the spectrum, extract the constant term of the heat trace, and the hole count falls out — the exact quantity excision writes.

5.1 The famous negative is a feature

Kac’s actual question — *can one hear the shape of a drum?* — was answered **no** in 1992, when Gordon, Webb and Wolpert constructed non-congruent planar domains with identical Dirichlet spectra. For thirty years that has been the punchline.

Here it is a security property. You cannot recover the geometry from the ringing. You *can* recover the payload, because the payload is χ and χ is one of the few things that survives into the spectrum. The medium supports a read that discloses the stored value **without disclosing the structure that holds it**. Nobody designed that; it is a consequence of what is and is not spectrally determined, and it costs nothing.

5.2 The fence, and it is load-bearing

Everything above is a two-dimensional planar-domain result. Excision as proposed is three-dimensional and the extension is *not* automatic. The scalar Laplacian does not reach the higher-dimensional parts of a space, so it would be surprising if its eigenvalues alone determined the Euler characteristic of a 3-manifold. The natural candidate is the **Hodge Laplacian**, and whether its spectrum determines χ is, as far as this pass found, open.

claim	status
planar domain: h holes $\Rightarrow a_0 = (1-h)/6$	THEOREM (Kac 1966; McKean-Singer 1967)
spectrum does not determine shape	THEOREM (Gordon-Webb-Wolpert 1992)
$b_k = \dim \ker \Delta_k$, per form degree	THEOREM (Hodge)
3D void network: scalar spectrum $= \chi$	NO — the scalar sector reaches b_0 only
unsectored spectrum $\Rightarrow$ Betti vector	OPEN
a physical probe that reaches degree k	THE ACTUAL QUESTION

The mathematics is settled and it settles in our favour. The Hodge theorem gives the k -th Betti number as the dimension of the kernel of the k -th Hodge Laplacian — the multiplicity of zero in that sector. The Laplacian determines the Euler characteristic. There is nothing to prove.

What is open is physical, and it is a better question. Three caveats survive, and together they reframe the whole experiment:

- 1. **You need the spectrum per form degree.** Whether the Betti vector can be read from an unsectored spectrum — not knowing which k -form sector each eigenvalue belongs to — is explicitly raised as unresolved in the discrete setting.
- 2. **Boundaries bite.** With a boundary the Hodge Laplacian’s kernel is infinite-dimensional; you need normal (relative) or tangential (absolute) boundary conditions to get finite kernels. A void network is a manifold *with* boundary, so this is not optional.

3. **Hodge isospectral examples exist** — Milnor’s counterexamples already were. That kills reconstructing the *shape* from the spectrum, which we do not want anyway (Section 5.1), and leaves the Betti numbers alone.

So the question stops being *is it possible* and becomes **which physical probe reaches which form degree** — which is answered by Section 5.3, and which is why rigidity turns out to be a mathematical choice rather than a comfort one.

And there is a cheap route around it if the answer is bad. A thin plate is effectively a 2D domain, so a *layered* architecture — holes written in stacked planar sheets — sits squarely inside the proven theorem. Less elegant than a solid block, fully supported by a proof, and it is how one would build the first demonstrator anyway.

5.3 Rigidity selects the operator, and the operator selects what you can hear

The substrate’s stiffness does not change its topology. **It changes the wave operator**, and the operator decides which Betti numbers are audible at all.

substrate	wave operator	what it reaches
membrane, gel, flexible sheet	scalar Laplacian	Kac applies <i>exactly</i> ; b_0
stiff plate	biharmonic, not Laplacian	the theorem must be re-derived
elastic solid	vector-valued elastic waves	the 1-form sector — where b_1 lives
gas- or vacuum-filled cavity	scalar acoustic only	b_0 , and nothing above it

Three consequences, and the first inverts the obvious intuition.

A flexible substrate is better for the bell read, not worse. A membrane under tension obeys the operator Kac’s theorem was proved for. **The cheap gelatin demonstrator of Section 4.2 is not just cheap — it is the right operator.** A stiff plate is the awkward case, because plate bending is biharmonic and inherits a different heat-trace expansion.

And the vector modes are the route to b_1 . A fluid cavity carries only scalar pressure modes, so it can never reach the 1-form sector no matter how well you ring it. A *solid* carries shear as well as compression — vector modes — which is the sector harmonic 1-forms live in. **Reading b_1 acoustically requires a solid, and that is a design constraint nobody would guess from the topology.**

The honest trade. Flexibility buys the right operator and costs channel purity: strain shifts frequency, so temperature, flexion and pressure all leak into the topology channel. Rigidity buys a stable channel and the wrong operator. A flexible substrate also absorbs kinetic damage and returns to shape without rupture — which protects the payload — while making the read more environment-sensitive. Neither is strictly better; **the matrix in Figure 5 is how you choose.**

5.4 We tried the numerical version and it failed

We attempted to recover h numerically from a computed heat trace on a finite-difference grid, twice. Both attempts failed: the discretisation puts an $O(1)$ offset in the constant term that does not cancel between domains, and a staircase boundary invents thousands of right-angle corners, each contributing exactly $1/16$ to the very term being measured. Switching to axis-aligned square holes did not rescue it.

The laptop version of the Kac read is real numerical analysis, needing a proper finite-element or spectral solver on the true curved domain. That is listed as work in Section 9, not as something already done.


*(A second route exists and is worth naming: model the **material** rather than the hole. The complement’s homology is determined by the object’s — Alexander duality — so the same numbers are reachable from either phase. It does not fix the discretisation error, which lives in the mesh and not in the choice of phase, but it is an independent check when the solver is right.)*

6. THE ALPHABET IS A PRODUCT

Three encoding layers exist as working reference implementations. Their sizes are not three design decisions:

PHYSICAL	eigenmode dictionary	3×3 grid: 9 modes + 3 chiral pairs = 12
CHANNEL	chiral winding codec	$m \in [-4,4] \times n \in [1,4] = 9 \times 4 = 36$
SEMANTIC	resonant codex	6 verbs × 6 subjects × 6 mod × 2 hands = 432
		12 × 36 = 432

Figure 4 · The alphabet is a product, not three inventions



The semantic alphabet is the product of the physical and channel alphabets. The codex size was never a design choice — it was forced by the grid and the cavity. That is Law 1 (the alphabet is the symmetry group) propagating up the stack.

Figure 4. The alphabet is a product. The semantic codex size is forced by the physical and channel alphabets, not chosen.

The semantic alphabet is the product of the physical and channel alphabets. The codex size was never chosen; it was forced by the grid and the cavity. That is Law 1 — *the alphabet is the symmetry group* — propagating up the stack, and it is the strongest single piece of internal evidence that the three layers are one object rather than three inventions.

A caveat found while checking: on a 3×3 square grid the eigenvalue $\lambda(p,q) = 4 - 2\cos(\pi p/N) - 2\cos(\pi q/N)$ gives $\lambda = 4$ with multiplicity three — the chiral pair (1,3)/(3,1) shares its eigenvalue with the *symmetric* mode (2,2). One of the three chirality channels is therefore spectrally inseparable from a non-chiral mode. A slight aspect ratio breaks the accidental degeneracy; the design target is the narrow window where the three-fold accident splits but the transpose pairs stay two-fold. **Better found now than after somebody builds a cavity.**

6.1 Chladni figures, and what they actually buy

Chladni patterns are the nodal lines of Laplacian eigenmodes. That identification is exact rather than analogical — and it is not decoration, because **it buys three specific things this architecture needs.**

A free visual verifier. The physical alphabet of Figure 4 is a set of eigenmodes, and a Chladni plate makes an eigenmode *visible in sand*. You can confirm a symbol is the symbol you meant by looking at it — no instrument, no calibration, no software. For a codebook that has to be agreed between two parties who cannot exchange a reference, that is not a small thing.

A shared reference nobody has to trust. Two readers with the same plate geometry see the same nodal pattern, because the pattern is a property of the boundary rather than of either reader. That is the same argument this paper makes for integer counts over calibrated thresholds, one layer down.

And a teaching surface. The three-layer stack is abstract until somebody watches sand jump into a pattern, and then it is obvious. A demonstration costing a speaker and a spoonful of sand is worth more to this architecture’s adoption than another page of algebra.

7. THE LOOP IS THE ADDRESS

The inversion that reorganises everything above.

Conventionally, you search a substrate for an invariant and walk loops to detect changes in it. Read-time holonomy. The instrument hunts.

Here, the path that writes is the loop, and the invariant sits where the path puts it. **Write-time holonomy.** You stop hunting and start placing.

And the shape of that is worth naming, because it is the whole architecture in one image. The invariant is deposited *on* a loop, and the invariant *is* a loop. The address closes and the datum closes, and they are the same walk. **Loop-the-loop** — the name Coney Island gave in 1901 to a track that returns you, upside down, to exactly where you started, which is what a holonomy does to a frame. The ride was the first of its kind, it terrified people, and the engineering question then was the same one as here: *does the thing come back around intact?*

Concretely: you install a **void graph**. Its topology is the invariant — that is what survives. The ordered walk through the network is the sentence. The abelianisation — which voids, how many, order forgotten — is the address. A pictogram is a subgraph with a chosen topology. Reading is traversing.

And the winding is the physically unclonable function. You cannot copy a topology by measuring an amplitude — a copier that reads magnitudes gets a shape, while the payload is a class. The disorder frozen in during writing is set by physical randomness and is not reproducible, including by whoever wrote it, so a challenge (an addressing pattern) yields a response governed by local disorder that is stable per device and uncorrelated across devices. The key is derived on demand rather than stored, so a powered-down device holds no extractable secret.

Tamper-response comes for free here rather than bolted on. Reading incorrectly — at the energies an invasive probe or a die-imaging attempt requires — is a supra-barrier operation, so **the attack anneals the very structure it is trying to read**. Conventional tamper-evident hardware detects intrusion and trips an alarm that can be defeated; here the destruction is a direct physical consequence of the read mechanism.

And the negative claim, stated explicitly because this is routinely oversold: this is not an unbreakable cipher. Unbreakability is a property of a protocol, never of a medium. The honest claim is an unclonable, tamper-responsive substrate for key storage and hardware authentication — which must still be designed against machine-learning modelling attacks, side channels (an acoustic read is a literal emanation), and thermal spoofing, since temperature shifts the Arrhenius behaviour an attacker could exploit. **This is never a cryptographic claim about the data itself.**

One honest note on tower addressing, and a clarification worth making explicitly because the idea invites a misreading. The tower is **bookkeeping, not geometry**. Nobody carves a trunk-and-branch superhighway through the glass; the construction is an *addressing scheme* for deciding which void is which, and the physical layout is free to be whatever the write head finds convenient. Mapping the address space onto the integer-forest construction buys exactly one real thing for free: arithmetic relations become spatial ones, so divisors sit near multiples and related data is physically close, unique and collision-free. It only pays if the data's natural relations *are* arithmetic. We hold it as an option, not a commitment. And the standing fence is unchanged: **this is never a cryptographic claim**. A faithful coordinate is the opposite of a one-way function.

8. SUBSTRATE FIDELITY — THE THREE LEDGERS

Every architecture here needs the same three-layer bookkeeping, and naming it once beats naming it six times.

The fidelity map. Per region: *how many distinguishable classes does this location still resolve?* Not “is it broken” — how derated is it. A region that once resolved a genus-8 handlebody and now resolves genus-3 is not failed; the encoder routes around it.

The local ledger. Each node's own parity or class. Cheap, fast, no coordination.

The universal ledger. The product around a global cycle. Because the structure group is abelian, **this is independent of the order you visit the nodes in** — which is why the gate matters for distributed systems and not just for crystallography. A non-abelian choice here would force a canonical ordering on every read, which in a drone swarm is a coordination protocol you cannot afford.

8.1 Two codes, opposite enemies — and this is the fence on flexible substrates

The single most useful thing this program has measured about noise is that **topological and magnitude codes fail to opposite perturbations**, and for a memory that maps directly onto two kinds of substrate damage.

- **Local damage** — a crack, a scratch, a defect, an inclusion — is a perturbation that does *not* wind around the loop.
- **Global deformation** — strain, thermal expansion, a bend — is a smooth turn applied to the whole medium, which *does*.

Measured in the reference codec, at **equal mean power** and with the decoder clamped to its codebook (symbol error rate, 9-symbol alphabet, chance = 0.889):

	local damage only		smooth global deformation	
σ	winding	magnitude	winding	magnitude
0.0	0.000	0.000	0.899	0.000
0.4	0.001	0.305	0.877	0.306
0.6	0.063	0.643	0.865	0.627
0.8	0.309	0.810	0.768	0.814
1.0	0.561	0.862	0.737	0.871

Read the two halves against each other. Under **local damage** the winding code makes **ten times fewer errors** at $\sigma = 0.6$ — 0.063 against 0.643 — which is the property a stored invariant is supposed to have. Under a **smooth global deformation** it inverts completely: at zero local noise the winding code is at **chance** while the magnitude code is **perfect**.

So the design consequence is sharp, and it is the honest fence on Section 5.3’s flexible substrates. Flexing a substrate is applying winding-scale common mode. A gel that absorbs a knock and returns to shape has protected the material and, in the same motion, forged the very turn the topological readout counts. **Rigidity is not only about which wave operator you get — it is about which of your two enemies you are choosing.** Glass resists the global deformation and cannot heal the local crack; gel heals the crack and invites the global deformation.

A correction we owe on this number. The codec as first written sent the winding symbol at amplitude 1 while the magnitude constellation averaged 0.391 in power — a **2.56× (4.1 dB) energy advantage** to the winding code — and its decoder returned unbounded integers rather than clamping to the alphabet. Both are fixed in the shipped tool, with the fixes commented in place. The table above is the corrected comparison. The law survives; the margin does not survive intact, and the pre-fix numbers should not be quoted.

8.1b The bell gives you two channels in one measurement

Resonance provides three genuinely independent channels:

Frequency — the topology channel. The heat-trace constant term. Carries χ .

Ring-down time (Q) — the loss channel. Set by damping, not topology. **Q is orthogonal to the payload**, so it measures the medium’s condition without touching the stored value. That is the fidelity map, physically realised: ring the block, and the decay rate tells you how healthy the substrate is while the frequencies tell you what it holds.

Mode ordering — the address channel. Set by cavity geometry. Neighbouring cavities are neighbouring bells.

And it self-frames. A ring is an event with a natural duration, so message length is set by physics rather than protocol. You need no framing bit and no stop symbol — **the ring-down is the frame**.

8.2 Every read is hot — and the barrier is a design parameter, not a constant

The coupling contrast that makes selective read possible is the same mechanism that deposits thermal dose into the payload. This should become the spec nobody publishes: **reads-until-erasure**, computable from ΔE , per-read dose, and $\tau = \tau_0 \exp(\Delta E/kT)$. Flash quotes *write* endurance. **A read budget is an unreported category** and this architecture should report one.

And that same exponential is a tier knob. Retention is set by thermally activated escape over whatever barrier protects the invariant, and lifetime is *exponential* in barrier height. With a typical attempt time $\tau_0 \approx 10^{-9}$ s at room temperature, a shallow barrier of $\Delta E \approx 0.5$ eV gives a lifetime of fractions of a second, while a deep barrier of $\Delta E \approx 1.1$ eV gives centuries. Two populations differing by a few tenths of an eV differ in retention by **many orders of magnitude**.

That yields something an archive-only medium cannot offer: **a continuum of tiers in one substrate**. A shallow tier is a self-erasing scratchpad whose fade is a *usable timer* — a command buffer that returns to a known-blank state on a predictable clock — beside an archival tier in the same material, with refresh-driven consolidation available in between.

The binning wall applies in the time domain too. Tiers must be separated by roughly ≥ 0.25 eV or their lifetime distributions overlap under realistic disorder, and the medium stops being able to tell a scratchpad from an archive.

(For excision the barrier is viscous flow of the network rather than defect escape, so the framework transfers and the numbers do not. Quoting them here is to show the shape of the knob, not to claim values we have measured.)

8.3 The read is separable from the record — and one read needs no pose

Because the payload is geometric, **storage and reading are separable concerns**. The record lives in the structure; the reader is an interchangeable probe. Light reads optical density and chirality, X-ray and electron probes read to far finer resolution, acoustic resonance reads the elastic state. **Storage density is therefore set by the write process, not by any one reader** — so a record written today can be read at higher resolution by a better probe tomorrow, and reading with several probes recovers more total information than any single one.

This has a consequence for the two read heads in Section 5 that is worth stating plainly, because it is an argument in the bell's favour we have not otherwise made.

A kinematic read requires registration. Dragging a probe across a record means specifying its rigid-body pose — three position and three orientation parameters — plus a start phase, seven numbers that lock probe to record. Without that reference the read clock drifts off-sample. Tomography inherits a version of the same burden.

A spectral read requires none of them. Ringing the whole object and reading the heat-trace constant term is pose-free by construction: there is no traverse, no start phase, no registration, and nothing to drift. The trade is that you get one number for the whole object rather than a spatial map — which is exactly the right trade for a *tier-level* or *page-level* integrity check, and the wrong one for random access.

8.4 Manufacturability, and why it is not a fantasy

The write head rides existing industry. Femtosecond direct-writing stations are commercial. Fused silica is a commodity. Acoustic input and output ride the surface- and bulk-acoustic-wave filter industry present in every handset. And self-organised void arrays mean the medium can pattern part of itself rather than requiring per-void placement.

The trade is explicit and should be stated rather than glossed: self-organised growth is cheap and supplies disorder for free, but its non-deterministic placement requires an **adaptive, learned addressing map** rather than fixed coordinates. You get the medium cheaply and pay for it in the address book.

8.5 What eats the breadcrumbs

Hansel and Gretel is the named failure mode, and it is precise: the birds ate the *path*, not the forest. Medium fine, readers fine, invariant gone. For excision specifically:

- **Viscous flow.** The only way to close a void is to move the network itself. Needs measuring, not asserting.
- **Crack propagation connecting voids.** The most likely real failure — and it is **readable**: b_0 drops, b_1 may rise. The medium reports its own damage in the units it stores.
- **Devitrification.** Affects everyone in this space equally.
- **The shell.** Every void carries a 20%-densified shock shell. Whether the shell relaxes faster than the void closes is unknown to us and could be the actual retention limit.

9. WHAT BREAKS IT

Five experiments, each with a bar to commit before the run. Two need a desk. Three need equipment we do not have, and those are written as a handoff.

Comparisons are to encoding classes and to theorems, never to products. That is a rule, and it also makes every test independently checkable.

DESK · E1 — the Hodge question

Does the spectrum of the Hodge Laplacian determine χ for a 3D void network? Literature plus mathematics. **Free, and it decides layered-plate versus solid-block before anyone spends anything.**

DESK · E2 — the Kac read, done properly

A finite-element or spectral eigensolver on the true curved domain, not a pixelated one. Recover h , then push past what any physical experiment could reach: add synthetic damping, truncate the spectrum, add measurement noise, and map where the read breaks. **That map is the competence surface**, and it is built entirely in software. *Bar*: recovered h equals the h built, with no free parameter beyond area and perimeter.

LAB · E3 — controlled connectivity in the archival substrate

Attempt genus-1, -2 and -3 void networks directly in fused silica without etching; verify by tomography. *Bar*: written genus equals intended genus. *If it fails*: layered-plate only. Not fatal.

LAB · E4 — Arrhenius retention, both encodings on one sample

Into one block write both a void pattern with known Betti numbers and a Type II birefringent pattern. Ladder the temperature by the standard accelerated-ageing procedure. Read both at every step. *Bar*: **the birefringence decays measurably while every Betti number stays an unchanged integer, up to the first topological event. That single plot is the thesis.** *If it fails* — if voids shrink continuously toward ambiguity rather than closing discretely — the class-encoding claim is dead and we say so.

The interesting outcome: if b_2 outlives b_1 , cavities outlasting tunnels, that is a **retention hierarchy** — the address layer more durable than the sentence layer. Worth reporting either way.

LAB · E5 — blind reader agreement

One party writes and records ground truth; a second reads without knowing it, computing Betti numbers with an established package rather than ours. *Bar*: 100% agreement on integer counts. Anything less means a threshold is hiding in the pipeline, and locating it is the result.

The cheapest physical test, if you have a transducer ring

Focus ultrasound into a gelatin phantom until cavitation, then **keep the phantom** — do not wash it. Image what remains after setting. *Bar*: a persistent, imageable void at the focus. Phantom gels only; never living tissue.

10. WHAT WE ARE NOT TESTING, AND WHY SAYING SO MATTERS

- **Not density.** We lose. Stating it first removes the only cheap shot available to a hostile reader.
- **Not write speed — but see Section 4.3, because “slow” without a number is useless and we owe you one.**
- **Not against anyone’s published lifetime figure.** Their number came from their oven and their samples; ours would come from ours. Comparing extrapolated numbers across laboratories is not a measurement.

11. PRIOR ART, AND WHAT SURVIVES IT

Three searches were run to decide whether this is a paper or a note. The result is the best possible shape for a proposal and the worst possible shape for a novelty claim: **every component exists, off the shelf, and nobody has assembled them.**

Established, cited, not claimed. Betti numbers of pore and void networks as a standard descriptor in porous-media modelling, cosmology and materials science. Femtosecond nanovoid formation and microexplosion physics in silica. Type III void gratings and their thermal limits. Femtosecond-written hollow microchannels of arbitrary configuration, including joined channels, 37 nm single voids, and nanofluidic channels formed by connecting voids. Kac 1966 and McKean–Singer 1967. Gordon–Webb–Wolpert 1992. Persistent homology and its software. Birefringent nanograting storage in fused silica. Chladni figures. Two-photon polymerisation and multi-beam interference lithography.

Not located. Storage of information in the *topology* of a void network rather than in void presence or absence. Persistent homology applied to laser-written structures. Controlled-genus writing as an encoding choice. The cavity/tunnel split as a deliberate selection of readout algebra at write time.

What survives, in one sentence:

The Betti numbers of a deliberately written void network are a payload, and because they are integers they fail discretely where a magnitude-encoded bit fails gradually.

That is an **assembly claim, not an invention claim**. It is weaker than we hoped and much more likely to survive review.

The obviousness risk, named. Because microfluidics builds these structures routinely, “store data in the topology” invites the reply that it is an obvious relabelling. The defence is not cleverness — it is E4. Nobody has argued, let alone measured, that a topological readout degrades discretely while a birefringence readout on the same sample fades. Measure that and the claim rests on evidence rather than on framing.

Prior-art status. No credentialed database pass has been run. The search covered publisher pages, arXiv and repository records. Novelty language is restricted to **not located**, never *new*.

12. WHAT THIS IS

Not a product. Not a density claim. Not a competitor to anything.

A specification of what a class-encoded memory would be, an algebra for choosing its readout at write time, a read head borrowed from spectral geometry, and five experiments that would end it — offered so that somebody with a fab, or an eigensolver, or a packet of gelatin, can decide whether it is worth their afternoon.

Everything load-bearing beneath it was discovered and named by someone else. The contribution is the assembly and the honest map of where it breaks.

DEPOSIT METADATA

Author Schulz, Matthew · Independent Researcher · Belmont, CA ORCID 0009-0000-0956-5743 · **License** CC BY 4.0 · **Type** Preprint

Related deposits. *The Twelve Answers Problem* — 10.5281/zenodo.21845003 · *Loop Walking: A Field Guide* — 10.5281/zenodo.21845063 · *Taking a Loop Walk Through the Integer Forest* — 10.5281/zenodo.21449570 · *The Tower Is Coordinate-Free* — 10.5281/zenodo.21798311 · *Walking the Loop: A Quick Guide to Holonomy* — 10.5281/zenodo.21448987 · *Reading the Sky Through Loops* — 10.5281/zenodo.21331652 · *The Loopy Knot Walker, Take Two* — 10.5281/zenodo.21817527

Reproducibility. The closed-word counts of Section 3.1 and the eigenvalue degeneracy of Section 6 regenerate from the scripts in the companion crate. The figures regenerate from `figures_p2.py`. Nothing else in this paper is a measurement by us, and nothing is presented as one.

Status. Speculative synthesis, fenced throughout. Every claim is tagged established, not located, doubtful, or open. Two of our own attempts at the numerical read failed and are reported in Section 5.3 rather than omitted.