

A Topological-Winding Memory

Signed Multilevel Storage in Chiral Spin Textures, with Upgradeable Multimodal Readout, Intrinsic Radiation-Hardness, and Tamper-Response

Working Draft v0.1 · Engineering / Proposal · Preprint, not peer reviewed · 2026

Matthew “Bear” Schulz

Independent Researcher

ORCID: 0009-0000-0956-5743 · mattschulz97@protonmail.com

License: Creative Commons Attribution 4.0 International (CC BY 4.0)

Companion to “Turtles, or Crystals” (Schulz, 2026), Zenodo DOI 10.5281/zenodo.20484438, and to the artificial-life paper of this deposit. This is a design proposal and analysis, not a demonstrated device.

Signed multilevel storage in chiral spin textures, with upgradeable multimodal readout, intrinsic radiation hardness, and tamper-responsive security

Preprint — not peer-reviewed. Companion engineering paper to the “Turtles, or Crystals” framework. CC BY 4.0. Draft v0.1 — prose for iteration. All references verified against primary sources.

Scope statement. This is a **design proposal and analysis, not a demonstrated device**. Every result here is analytic (closed-form bounds) or simulation (information-layer models); **no physical cell has been fabricated or measured**, and the central physics question — whether crystal fracture quantizes the topological linking number (§8, the make-or-break) — is identified, not yet run (it requires micromagnetic simulation on hardware we did not have). The contribution is the architecture and the falsifiable test set, stated so each claim can be killed. Read it as a roadmap with numbers.

Abstract

Charge-based memory stores a bit as electrons on a gate, a representation that is volatile, radiation-soft, and read-destructively coupled to the medium that holds it. We describe a memory in which the stored quantity is instead a **topological winding** — the signed, quantized invariant of a chiral spin texture — and show that this single change in representation propagates into four concrete advantages. First, the winding is a *signed, multilevel* symbol: sign encodes chirality and is read for free by circular dichroism, while magnitude encodes the charge of a skyrmion “bag,” giving a balanced, ternary-friendly alphabet whose capacity we bound. Second, retention is governed by a single Arrhenius barrier, so one material yields a *continuum of memory tiers* — a self-erasing command buffer and a years-stable archive — by barrier height alone. Third, because the bit is geometric rather than electronic, the **stored record is decoupled from the reader**: it can be interrogated by any of several physical probes (optical, X-ray, RF, acoustic), so storage density is fixed by the write process and reads remain upgradeable indefinitely. Fourth, the same physics that protects the bit makes the device **intrinsically radiation-hard** and enables **tamper-responsive security**, in which an invasive read physically anneals the secret it is reaching for. We give the governing relations, simulation results locating the capacity and over-binning limits, a screen of real candidate substrates with the specific engineering each would require, an honest benchmark against flash/DRAM/MRAM/ReRAM, and a set of falsifiable predictions with explicit conditions under which each claim fails. The problems this targets are concrete: radiation-hard non-volatile storage for environments that destroy charge memory; archival fidelity that survives reader obsolescence by decoupling the record from the reader; higher per-cell density via a signed multi-axis alphabet; and security plus self-timed erasure built into the physics. We do not claim to beat DRAM on speed or flash on near-term density, and we say where we lose as plainly as where we win.

In plain words (the device explained to a smart 12-year-old). Today a computer stores a bit as a tiny pile of electric charge — like a little puddle of water. Puddles leak (the data fades), a splash knocks them over (radiation flips the bit), and to check how full a puddle is you have to dip a finger in and disturb it (reading wears it out). This memory stores the bit a completely different way: as a tiny, stable **twist** in a magnetic material — a microscopic whirlpool called a *skyrmion*. A twist is a whole number; you can't half-untwist it by nudging, only by paying a real energy “toll,” so it doesn't leak and a radiation splash can't flip it. You can count *higher* than 0/1 by nesting whirlpools inside whirlpools, and the twist can spin left or right — a built-in plus or minus sign. You can read it from a distance with light, sound, or X-rays without touching it, and you can invent better readers later without rewriting the data. And if a spy tries to read it the hard way — by heating it up to image it under a microscope — that very act smooths the twist away and erases the secret. *A full step-by-step build, in the same plain language, is the companion*

1. Introduction

The dominant memory technologies store information as charge. A flash cell holds electrons on a floating gate; a DRAM cell holds them on a capacitor that must be refreshed. This representation is responsible for several of the field's hardest constraints at once: charge leaks (limiting retention), a stray ionizing particle can add or remove it (a single-event upset, the root of radiation softness), and the act of reading is electrically coupled to the act of storing.

We take a different representation. The stored quantity is a **topological invariant** of an ordered medium — concretely, the winding number of a chiral spin texture such as a magnetic skyrmion. Topological invariants have a defining property useful to a memory engineer: they cannot be changed by any continuous local deformation, only by an event that crosses an energy barrier. Particles realized as field windings are a sixty-year-old idea in physics [Skyrme 1962], and their classification by the homotopy of the order-parameter space is standard [Mermin 1979]; magnetic skyrmions as storage elements are an active engineering field. Our contribution is not the skyrmion-as-bit, which is established, but a specific *architecture* built on it: a signed multilevel alphabet read by independent physical channels, a single-material multi-tier retention scheme, a reader-independent ("gramophone") storage model, and a security model based on tamper-response. We state each as a quantitative, falsifiable claim.

A note on how to read this paper. It deliberately crosses several fields that are not usually held in one frame — coding and information theory (the alphabet and its capacity), condensed-matter defect physics (the winding itself), optics (the orbital-angular-momentum sibling), thermodynamics (retention), hardware security (the PUF), and acoustics (the read). Each jump between them is a load-bearing step in the argument, not an aside, so at the hardest crossings we include a short plain-language note in a quoted box. Those notes are not a simplification for non-experts; they are a hinge for *experts* — a specialist deep in one of these fields who has never had reason to enter another. The argument is the stitching of these domains together; the boxes make the stitches visible.

A statement of method and stance. We should be explicit about what this paper is and is not, because the fields it touches have earned the right to be addressed with care. This is a *synthesis*: a scaffold built by combining established results across disciplines in a way we believe is new, offered in earnest. We are students of the literature, not veterans of any one of these benches — and our contribution is precisely the cross-disciplinary architecture, not a verdict on any single domain's physics. Where we propose a materials route or an engineering recipe, read it as a well-intentioned observation from an outside vantage, not as an instruction to specialists who know their field far better than we do. The novelty is the reason we publish at all: a synthesis that proposes nothing new would not be worth a reader's time, so we take the synthetic risk on purpose. But a bold synthesis is only as good as the testing it invites. So we hand the reviewer the hammer and ask them to test and strengthen this scaffold through its academic destruction: what survives demonstrates the merit of the idea, and what is destroyed is offered a second life through the opportunity to engineer a new solution in its place. The companion *Seam & Assumptions Register* names, discipline by discipline, where we most need that hammer to fall.

2. The cell: a signed multilevel alphabet

The natural symbol is the topological charge Q of a texture in one cell. Three states are robust and physically distinct: an antiskyrmion ($Q = -1$), the uniform background ($Q = 0$), and a skyrmion ($Q = +1$). This is a **balanced ternary** cell, and the choice of radix is principled rather than aesthetic: the radix economy of base-3 is the closest integer to the optimum e , and the three states map to the three most distinguishable physical configurations. The sign is the chirality of the texture and is read directly by circular dichroism — opposite windings absorb left- and right-circularly-polarized light differently — so the sign bit requires no decoding.

In plain words: instead of a switch that is only off/on (0/1), each cell is a little whirlpool that can spin one way, the other way, or not at all — three values, with the spin direction acting as a built-in plus-or-minus sign. Counting in threes (not twos) is, perhaps surprisingly, close to the most efficient way to count at all.

Magnitude extends the alphabet. A skyrmion "bag" nests N textures within an outer wrapper and carries total charge $Q = N - 1$, a result demonstrated experimentally with charge tunable to ~ 10 and stable at room temperature [Foster et al. 2019; subsequent room-temperature bag work]. A cell therefore carries a *signed integer*, and a multi-axis cell — sign, magnitude, and one geometric (orientation) axis — multiplies independent channels. Here the argument makes a deliberate cross-domain jump worth flagging, because the payoff is a property the whole scheme depends on: *the same winding idea exists in light, not just in spin*. A vortex beam carries quantized **orbital angular momentum** $\ell\hbar$ per photon, where ℓ is the winding number of the optical phase — angular momentum with no rest mass, the optical sibling of the spin-texture winding [Allen et al. 1992]. Its relevance here is twofold: it is an additional quantized, field-read symbol axis, and — the reason it earns a place in a memory paper — it makes explicit a property the whole scheme relies on, that *the topological charge is independent of the quantity of medium* (a quantized circulation does not

depend on how much fluid circulates), so the stored invariant does not scale with cell volume. The optical case is the cleanest illustration of a fact we need to hold for the spin case: winding is a count, not an amount.

How many dimensions pay? Capacity per channel follows Shannon, $\text{bits} \approx \frac{1}{2} \cdot \log_2(1 + \text{SNR})$; an added axis is worth keeping only while its SNR stays well above the distinguishability floor. A Monte Carlo encode→noisy-read→decode model (Fig. 1, panel B) shows net reliable bits per cell **peaking at four dimensions and declining beyond**, because additional axes share the same physics and arrive with SNR near unity, so their error-correction overhead exceeds the raw bits they add. This peak is a property of *correlated* axes — degrees of freedom in one cell that share physics and so suffer declining per-axis SNR; strictly independent, equal-SNR axes would not self-limit (a separate simulation confirms dimension grows unbounded when axes do not interfere). The four-dimension figure is therefore a statement about coupled channels in a real cell, not a universal dimension cap. The practical recommendation is **three robust dimensions (27 states) or four aggressive ones (81 states)**; five or more is bloat.

The hard wall is distinguishability. Reliable levels per axis $\approx$ dynamic range / ($\sim 6\sigma$ of read noise); packed tighter, adjacent level distributions overlap and the symbol-error rate climbs super-linearly past a critical density — the error catastrophe, in the spatial domain. Our simulation locates this cliff at a level count consistent with the 6σ rule (Fig. 1, panel A). **This is the engineering claim most exposed to falsification: if a fabricated cell cannot hold ~ 27 distinguishable states at the predicted read SNR, the density claim fails.**

3. Two-tier retention from one material

Retention is set by thermally activated escape over the barrier that protects the winding: $\tau = \tau_0 \cdot \exp(\Delta E / k_B T)$. The exponential dependence on barrier height is the lever. With a typical attempt time $\tau_0 \approx 10^{-9}$ s at room temperature, a shallow barrier ($\Delta E \approx 0.5$ eV) gives a lifetime of fractions of a second — a genuinely volatile **command buffer** that returns to a known-blank state on a predictable clock — while a deep barrier ($\Delta E \approx 1.1$ eV) gives centuries (Fig. 1, panel C: our model spans ~ 0.5 s to ~ 400 years across this range). Because lifetime is exponential in ΔE , two populations differing by a few tenths of an eV differ in retention by many orders of magnitude.

Multiple barrier heights in one substrate are obtained by defect species, by dopant pinning (a Cottrell atmosphere raises the barrier), by patterned anisotropy, or by texture size. The result is a **continuum of tiers in a single material**: a self-erasing scratchpad whose fade is a usable timer, and an archival tier, with the option of refresh-driven consolidation in between. Tiers must be separated by $\geq \sim 0.25$ eV so their lifetime distributions do not overlap given realistic disorder — the binning wall, now in the time domain.

4. Addressing and readout: the gramophone decoupling

Because the bit is geometric, *storage* and *reading* are separable concerns, and we exploit this explicitly. The record lives in the three-dimensional texture geometry; the reader is an interchangeable probe dragged across it. Light reads optical density and chirality; X-ray and electron probes read to ångström and atomic resolution; RF and microwave read the spin state; an 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 the same record with several probes recovers more total information than any single one.

In plain words: a vinyl record holds the music in its grooves; the needle is just the thing that plays it, and you can build a better needle later without re-cutting the record. Here the data is in the shape, and the reader — light, sound, X-ray — is the swappable needle.

For write and fine addressing without an optical window — the relevant case for embedded, sealed, or harsh-environment use — we pair a microwave field (which couples to spin, performing the winding write) with a gigahertz acoustic standing wave (whose ~ 0.6 μm wavelength localizes to the feature scale and reads by strain-induced frequency shift). The two channels are independent because they couple to different degrees of freedom; their intersection is a contactless address. Different tiers and species resonate at different frequencies, so tier selection is achieved by frequency rather than by spatial wiring (Fig. 2).

The read is kinematic: a probe is moved relative to the record, and only the *relative* motion matters (the record may translate past a fixed probe or the reverse). This has three consequences worth stating. First, **the read supplies a dimension the record need not store**: a static record is inert, and the signal exists only as the probe traverses it — a three-dimensional structure is recovered as a time-ordered series of two-dimensional reads, so reconstruction fidelity scales with the number of slices traversed (the tomographic exchange rate). Second, the record may store a *generating function* rather than raw samples — coefficients the read evaluates along its path — which trades storage for read-time computation and, for a lawful (compressible) texture, recovers a high-resolution field from a small payload. Third, a kinematic read requires **registration**: specifying the read fully needs the probe's rigid-body pose (six degrees of freedom, three position and three orientation) plus a start phase (where in a revolution the read begins) — seven parameters that lock the probe to the record. Absent this reference, the read clock drifts off-sample, the same failure

mode as an uncorrected copy.

A read-velocity profile must also be chosen, and the choice is the same one made by rotating optical and magnetic media. Constant *angular* velocity (a fixed RPM) gives a simple mechanism but a position-dependent data density, since linear velocity $v = \omega r$ falls toward the center; constant *linear* velocity holds density uniform at the cost of a variable drive. This is a genuine design parameter, not a metaphor. Finally, two failure modes familiar from contact media — wear and physical damage — are addressed at the level of representation rather than handling: a **non-contact field read removes the wear term entirely** (there is no stylus to foul the record or to be fouled), and because the stored quantity is a topological invariant, **local geometric damage such as a scratch or a warp cannot change it** — only a barrier-crossing event can. The medium's fragility, in other words, is not in the same place as its data.


4a. Coarse reading as a classifier: the low-pass triage

A multilevel record admits a read strategy that a binary one does not, and it is worth showing in full because it inverts the usual instinct that a read should always be as sharp as possible. Deliberately read the record at *low resolution* first — coarsen the amplitude bins, blur the spatial detail — and the coarse pass acts as a **classifier**: it answers "what kind of record is this" cheaply, before any expensive high-resolution read answers "which exact value." The architecture then pays for a fine read only on the candidates that survive triage.

In plain words: you can recognize a friend's silhouette through fog before you can make out their face. Blurring throws away the fine detail that distinguishes one face from another but keeps the coarse shape that says "person, not lamppost." A cheap blurry look sorts the record into a category; only then do you pay for the sharp look.

The mechanism is general and not specific to this medium: fine detail is where individual records differ, coarse structure is where *categories* live. Discarding the high-spatial-frequency content removes the within-class variation while preserving the between-class signal, so for the category question the coarse read's signal-to-noise is actually *higher* than the full read's. We measured this directly on the stored alphabet: a deliberately lossy, coarse-binned read classified category as well as or better than a full high-dimensional read at roughly **sixty-fold compression** (Fig. 1e). There is a window — blur too little and you have paid for detail the category question never used; blur too much and the categories themselves smear together and the classification fails.

That window is the point of contact with the over-binning limit of §2, and the two are the same knob seen from opposite sides. Over-binning is what happens when amplitude levels are packed *past* the point where the read noise can separate them — category-blur pushed until the categories collapse, an uncontrolled failure. Controlled coarse reading is the same coarse-graining used *deliberately* and stopped inside the window — a designed feature. One mechanism: suffered as a limit when it runs past the edge, exploited as a triage when it is held back from it. The idea is borrowed frankly from outside memory engineering — it is a coarse-graining argument of the kind familiar in information theory and the study of complex systems, applied here to a read — and it is the clearest single instance in this paper of a benefit that only appears when one looks at the device through more than one discipline's lens.



5. Intrinsic radiation hardness

The radiation softness of charge memory follows directly from its representation: a bit is a pile of charge, and an ionizing event adds or removes charge. A winding is not a pile of charge; it is a non-local invariant of a field. An ionizing strike deposits energy locally and broadband, whereas flipping a winding requires energy delivered to a specific resonant mode above the barrier. Local broadband deposition is not resonant supra-barrier deposition, so the generic event deposits heat the lattice dissipates without changing the invariant.

Two protections multiply rather than add. The first is topological: local damage cannot reach a non-local invariant. The second is the host: a wide-bandgap material such as silicon carbide (gap ~ 3.3 eV vs ~ 1.1 eV for silicon) requires far more energy to liberate carriers and recovers thermally. The combination is meaningfully harder than charge-based flash, which is the comparison that matters for space, reactor, and defense deployment.

6. Security: unclonable and tamper-responsive

The disorder frozen into the medium during growth — the exact defect and texture map — is set by uncontrollable physical randomness and cannot be reproduced, including by the original manufacturer. This is a physical unclonable function (PUF): a challenge (an addressing pattern) yields a response (the readout governed by local disorder) that is stable per device and uncorrelated across devices, and the key is derived on demand rather than stored, so a powered-down device holds no extractable secret.

The distinctive property is **tamper-response**, which our architecture provides natively. Reading incorrectly is a supra-barrier operation; an invasive probe or an attempt to image the die at the energies required for microscopy therefore **anneals the very textures being read**. Where conventional tamper-evident hardware detects intrusion and trips an alarm that can be defeated, here the attack destroys its target as a direct physical consequence of the read mechanism. Self-destructing and tamper-evident PUFs exist as separate primitives [Pappu et al. 2002; Gassend et al. 2002]; the contribution here is that destruction-on-invasive-read is intrinsic to the storage physics rather than a bolted-on circuit.

We make the negative claim explicitly, because it is routinely oversold: **this is not an unbreakable cipher**. Unbreakability is a property of a protocol (a one-time pad), never of a medium. The honest claim is an unclonable, tamper-responsive, self-erasing substrate for key storage and hardware authentication, which must still be designed against machine-learning modeling attacks, side channels (the acoustic read is a literal emanation), and thermal spoofing (temperature shifts the Arrhenius behavior an attacker could exploit).

7. Manufacturability

The stitch from the security section is the key one, because it turns what looks like a manufacturing *problem* into the manufacturing *route*. A PUF needs disorder that no one can reproduce; the cheapest way to get unreproducible disorder is to stop trying to place every bit precisely and instead let the medium organize itself as it grows. So the same self-organized growth that would be a defect in a lithographic process — you cannot say exactly where each texture lands — is precisely what supplies the security primitive for free. The representation does not merely tolerate cheap manufacturing; it is made *better* (more unclonable) by it. That is why the architecture rides existing industries rather than demanding a new one.

The architecture rides three existing industries. Silicon carbide is already produced at wafer scale by the power-electronics sector and supplies the wide bandgap, the radiation hardness, and hollow-core dislocations usable as access channels. Acoustic input/output is the surface- and bulk-acoustic-wave filter industry present in every handset. And self-organized texture formation means the medium can pattern itself rather than requiring per-bit lithography. The trade is explicit: self-organized growth is cheap and supplies the PUF disorder for free, but its non-deterministic placement requires an adaptive, learned addressing map rather than fixed coordinates.

8. Validation: simulations, and how each claim dies

We report a Monte Carlo information-layer simulation (Fig. 1) that confirms three quantitative predictions: the over-binning cliff at a level count consistent with the 6σ rule; the peak of net reliable capacity at four dimensions; and the multi-tier retention spread from sub-second to centuries across a plausible barrier range. The simulation validates the *information architecture*; it does not validate the spin physics, which requires micromagnetic simulation (mumax³ or OOMMF) and, ultimately, fabrication.

We state the falsification conditions plainly, because a claim that cannot be killed is not an engineering claim:

- **Density** dies if a fabricated cell cannot hold ~ 27 distinguishable states at the predicted read SNR, or if net capacity does not peak near three to four dimensions.
- **Retention** dies if measured lifetimes do not follow $\tau_0 \cdot \exp(\Delta E/k_{\text{BT}})$ across engineered barrier populations, or if tiers cannot be separated cleanly at ~ 0.25 eV spacing.
- **Upgradeable read** dies if independent probes cannot recover the same record above the noise floor, i.e. if the stored geometry is not in fact modality-independent.
- **Radiation hardness** dies if single-event upset rates are not materially below charge-based flash under matched fluence.
- **Tamper-response** dies if a below-threshold invasive read can extract the state without annealing it.

Beyond these self-tests, the companion **Seam & Assumptions Register** lists, for each load-bearing assumption, which domain expert should interrogate it, our honest confidence (separated from how assertive the prose reads), and the exact question we could not answer ourselves. Several of those assumptions — micromagnetic switching energies, bag-state retention, SiC integration — are calibration claims where confident prose is least trustworthy. If you are a specialist, that register is where you are most needed.

8a. The substrate problem: the wall this architecture stands or falls on

It would be dishonest to present the preceding sections as a design awaiting only fabrication, because there is one demand underneath all of them that is not yet met by any known material, and naming it plainly is the most useful thing this paper can do. **Every advantage above is contingent on finding a single substrate that satisfies several stringent properties at once** — and that conjunction, not any individual property, is the wall.

The architecture asks one material to simultaneously: (i) host **room-temperature-stable** chiral textures at zero or low field; (ii) support **tunable "bag" charge** so the alphabet is multilevel, not just ± 1 ; (iii) present **several well-separated retention barriers** ($\geq \sim 0.25$ eV apart) so the single-material multi-tier scheme works; (iv) couple cleanly to **at least two independent read channels** (e.g. optical chirality and acoustic strain); and (v) tolerate, ideally exploit, the **self-organized disorder** that supplies the security primitive. Each of these has been demonstrated *somewhere*, in some material, in isolation. No material is known to do all five together, and there is no guarantee the set is jointly satisfiable — the properties pull against one another (deep barriers fight room-temperature writability; tunable multilevel charge fights clean two-channel readout).

We state the parallel deliberately, because it is the honest frame and a cautionary one. **This is the same shape of problem that has dogged artificial life for fifty years.** ALife reasoned, correctly, that life is organization rather than matter — and then repeatedly failed to *instantiate* open-ended evolution, because the substrate it ran on quietly refused to supply something the organization needed (the companion paper reports exactly such a wall). The lesson transfers: an architecture can be substrate-independent in its *logic* and still be entirely hostage to substrate in its *realization*. Saying "here is a beautiful memory architecture; a materials scientist need only supply the medium" is the engineering cousin of "here is a beautiful definition of life; a chemist need only supply the cell." The supplying is the hard part. We do not want to repeat ALife's optimism with the materials step waved through.

So rather than assert the material exists, we offer the search itself as the contribution, with candidate directions a materials group could falsify quickly:

- **Multilayer DMI films (Pt/Co/Ta-class).** The most mature route to room-temperature skyrmions; strongest on (i), weakest on (iii) — engineering several clean retention tiers in a thin-film stack is unproven. *First test:* can two skyrmion populations with ≥ 0.25 eV barrier separation coexist and be independently addressed in one stack?
- **Van der Waals magnets (Fe_3GeTe_2 , CrI_3 and successors).** Attractive because the layered structure offers both clean interfaces and, via the companion paper's logic, clean cleavage; weakest on room-temperature stability for the early members. *First test:* room-temperature bag-charge tunability in a vdW monolayer or few-layer.
- **Bag charge as the multilevel axis (Foster-class liquid-crystal / ferromagnet systems).** The experimental anchor for tunable magnitude; the open question is whether bag states are retentive and readable, not merely creatable. *First test:* measured retention spectrum of bag states, not just their existence.
- **Silicon carbide as host, textures as guests.** SiC supplies (iv) and (v) and the radiation hardness outright; the question is whether a texture-bearing magnetic layer can be integrated with it without losing the wide-bandgap advantages. *First test:* a SiC-integrated magnetic layer that still upsets below flash under matched fluence.

The honest status, then: this is a proposal whose central risk is not any one of its five claims but the **existence of their intersection in one material**. The fastest way to kill or advance the whole program is to attack that intersection directly — and the value of laying the architecture out in full is precisely that it tells a materials scientist *which* conjunction to hunt for, and which single measured failure (no two-tier coexistence; bag states that will not hold; an upset rate no better than flash) would close the question.

8b. Falsifying the candidates now: a materials screen with engineering recipes

The substrate problem (§8a) is not a reason to stop; it is a reason to do the materials triage on the page, with the literature in hand, so the reader sees which existing options *already fail* and why, and what specific engineering would have to be true to rescue each. We screen the leading candidates against the five-property conjunction and, where a property is missing, propose the concrete intervention — stated as a falsifiable design conjecture, not a result. The point is to convert "find a magic material" into "here is the short list, here is exactly where each one breaks, here is the doping or stacking move that would fix it without breaking something else."

The screen (what each candidate buys you, and where it dies).

Candidate	Buys you	Fails at	Engineering rescue (conjecture, with the thing it must not break)
Pt/Co/Ta (and Pt/CoFeB/MgO) DMI multilayers	Room-temp skyrmions, CMOS-adjacent fab, electrical read	One barrier height only $\rightarrow$ no native retention tiers; placement is stochastic	Graded-anisotropy stack: vary the heavy-metal/ferromagnet interface thickness layer-to-layer to write ≥ 2 barrier populations ≥ 0.25 eV apart — <i>without</i> pushing the deepest barrier so high it can't be written at room temperature (the writability/retention tension of §3)
Van der Waals magnets (Fe_3GeTe_2)	Clean layered interfaces; clean cleavage (the bridge to the companion	Curie temperature below room T for	Interface proximity / pressure / electrostatic doping to lift T_c above 300 K — <i>without</i> destroying the layer cleanliness that made the cleavage and the sharp barriers attractive in the first place

(rescue, CrI ₃ , CrSBr...)	bridge to the companion paper); atomically sharp tiers	below room T for the early members	the cleavage and the sharp barriers attractive in the first place
Liquid-crystal / ferromagnet "bag" hosts (Foster-class)	The only demonstrated tunable multilevel charge ($Q \approx \pm 10$) — the multilevel alphabet's anchor	Retention of bag states unmeasured; readout-at-density unproven	Pin the bag with a dopant or anisotropy template that deepens the bag's escape barrier — <i>without</i> collapsing the level spacing that lets ~10 charges stay distinguishable (the over-binning wall of \$2)
Silicon carbide (host, textures as guests)	Rad-hardness and wide bandgap outright; hollow-core micropipes as read channels; wafer-scale supply	SiC itself is not the magnetic medium	Heterointegrate a thin texture-bearing magnetic layer on/in SiC — <i>without</i> introducing carrier-trap interface states that erase the wide-bandgap rad-hardness advantage

Each rescue row is a single experiment a materials group could run and report against us. And each "must not break" clause is the honest part: these properties are coupled, so every fix risks paying for one property with another, and the screen names that bill rather than hiding it. The fastest disproof of the whole architecture is a clean demonstration that one of these conjunctions is *not* jointly reachable — that lifting T_c always destroys cleavage, or that pinning bags always collapses their spacing. We would rather that be found quickly than assumed away.

The doping/engineering primitives we are actually leaning on (all established mechanisms, cited so the conjectures are anchored): dopant pinning via a Cottrell atmosphere to raise a local barrier; graded deposition to write a position-to-property map during growth; interface proximity coupling to raise ordering temperature; strain and electrostatic gating to tune anisotropy post-growth. None of these is new physics; the conjecture is only that they can be *combined* to hit the five-property target.

8c. What this buys against the incumbents: an honest benchmark

A new memory earns attention only by saying, in the incumbents' own units, what it does better, what it does worse, and which problem it solves that they cannot. We are explicit that **our column is a projection from the architecture, not a measured device**; the incumbent columns are real, sourced figures. The comparison is meant to be argued with.

Property	NAND flash	DRAM	STT-MRAM	ReRAM	This proposal (projected)
Non-volatile	yes	no	yes	yes	yes (and <i>tunably</i> so — see retention)
Endurance (cycles)	~10 ⁴	>10 ¹⁵	~10 ¹⁴ –10 ¹⁶	~10 ¹²	no wear term on read (topological, non-contact); write endurance set by host
Retention	~10 yr	ms (refresh)	~10 yr	~10 yr	~0.5 s → centuries, <i>selectable in one material by barrier height</i> (\$3)
Radiation hardness	soft (SEU-prone)	soft	hard	hard	hard×hard (non-local invariant <i>and</i> wide-bandgap host, \$5)
Bits per cell	1–4 (MLC/QLC)	1	1 (MLC hard)	1–~3	signed multilevel: ~27 (3 axes) to 81 (4), \$2
Read model	fixed sense	fixed sense	fixed sense	fixed sense	reader-decoupled: upgrade the probe, not the media (\$4)
Built-in security	none	none	none	none	PUF + tamper-response intrinsic (\$6)

Read this table honestly and it says four things. **(1) We do not win on speed or near-term density** — DRAM and MRAM are faster, and a shipping QLC flash die is denser today than a first-generation self-organized substrate would be. We are not a DRAM replacement and should not be sold as one. **(2) The wins are specific and they compound where they matter.** Rad-hardness is *doubled* — the bit is a non-local invariant a local strike can't reach, *and* the host is wide-bandgap — which is the axis that matters for space, reactor, and defense payloads, and it is the single most defensible claim in the paper. **(3) The architecture is genuinely smarter in two places competitors structurally cannot copy:** the signed-multilevel alphabet packs more distinguishable states per cell than a charge level can (because chirality is a free, orthogonal sign axis no charge memory has), and the gramophone decoupling means a record written once can be re-read at higher fidelity by a better instrument *decades later* — an archival property no fixed-sense memory offers. **(4) The retention continuum is a category nobody else has:** a single material that is simultaneously a self-erasing scratchpad and a century-scale archive, set by barrier height, replaces what is today two or three separate memory products.

So the problems we actually solve, stated plainly: **radiation-hard non-volatile storage for environments that destroy charge memory** (the lead application); **multi-decade-to-century archival fidelity that survives reader**

obsolescence (write it now, read it with tomorrow's probe); **higher per-cell information density via a signed, multi-axis alphabet** where the sign bit is free; and **storage with security and self-timed erasure built into the physics** rather than bolted on. The compression/density gain is real but modest and contingent (≈ 3 robust axes $\approx$ a 27-symbol alphabet, ~ 4 -5 bits/cell vs flash's 3-4, *if* the distinguishability holds); the rad-hardness and the reader-decoupled longevity are the claims we would lead with, because they are the ones the incumbents cannot reach by iterating.

9. Limitations

The microwave channel cannot perform fine spatial addressing (diffraction limit), which is why optical or acoustic probes set resolution. The capacity figures assume clean, fixed barriers; real disorder broadens every distribution and is the dominant source of read error. The security claims are primitives, not a cipher. And the self-organized manufacturing route trades placement determinism for cost and security, requiring adaptive addressing.

The over-binning limit (§2) is one instance of a more general pattern worth naming: a packing knob buys capacity cheaply until a critical density, past which loss turns non-linear. Over-binning is its *read-domain* form (amplitude levels packed inside the read-noise floor $\rightarrow$ error catastrophe); the same shape recurs as defect-density-driven structural collapse in the storage medium and, in the companion evolvability paper, as the mutation error threshold. These are shared phenomenology rather than one equation — the mechanisms differ — and they differ in recovery: over-binning is recoverable by reading a coarser (e.g. ternary) alphabet, whereas a structural collapse is not. The practical consequence for this paper is the same conservative rule the capacity claim already follows: operate back from the maximum, at a density the read noise can actually resolve.

Appendix A. Implementation primitives: the device as a topological circuit

A reader who has followed the physics of §§2-6 may still balk at the move this appendix makes, so it is worth stating the hinge before making it. *Why should a magnetic whirlpool have anything to do with an electrical circuit element?* The connection is not metaphor: a spin texture is a configuration of a material whose **electrical response** — its resistance, the voltage it induces when it moves, the fields it couples to — is set by the texture's geometry. Change the winding and you change the resistance; move the winding and you induce an EMF. So the texture is not *like* a circuit element; the texture's geometry *is* the value of a circuit element, read out electrically. That is the bridge from the spin-physics picture to the engineer's parts bin, and it is what lets the whole device be described in components a circuit designer already owns. The architecture above is therefore not an analogy to electronics; it is realizable in established spintronic/topological components, several of which perform more than one role at once. We grade each mapping **[REAL]** (established physics, has literature), **[ADJ]** (real mechanism; the circuit-element framing is a mapping), or note the speculative joint explicitly.

The memory cell is a memristor [REAL, keystone]. Domain-wall and skyrmion memristors — resistance set by texture position or winding count, persistent across power cycling — are an established, funded neuromorphic-device class. The signed-multilevel cell of §2 is a multilevel memristor; this places the substrate within existing hardware rather than adjacent to it.

The winding is a self-reading inductor [REAL]. A winding stores energy in its field (an inductor); a *moving* winding generates an emergent electric field (emergent electrodynamics — the topological cousin of Faraday induction). A texture in motion therefore produces its own readout signal: sensing and power co-located in one geometric object, with no external probe required for motion detection.

The defect is a triple-duty transducer [REAL]. A protected disclination simultaneously stores anneal-resistant state (memory), concentrates curvature (shapes the structure), and polarizes via the strain gradient at its core (flexoelectricity — measured atom-by-atom at dislocations in SrTiO_3). The same flaw that holds the bit converts mechanical events into an electrical signal — relevant to the tamper-response of §6, since a mechanical intrusion is transduced into a readable electrical trip.

Read/write hardware [REAL]. A spin-torque oscillator converts DC current into GHz precession with no external clock (self-sourced write/read carrier). Photonic/phononic crystal structures with line defects route the carrier (on-chip waveguide bus) and, as bounded periodic regions, define the size-set resonant cells and frequency-selective filters of the addressing scheme in §4. Backscatter (reflection modulated by winding state) supplies passive, transmitter-free readout.

Frequency-division addressing [REAL]. Cell size sets a cell's resonant frequency, so an array of differently-sized cells presents a set of independently addressable frequency channels. Cells whose resonant frequencies stand in commensurate (small-integer) ratios couple and tend to merge into a shared channel, while incommensurate ratios remain isolated — standard coupled-resonator behavior. Channel isolation is therefore set by frequency-ratio spacing, giving a direct design knob for crosstalk in the addressing scheme of §4.

The clock [REAL]. Timing is supplied by a self-sourced oscillator (e.g. a spin-torque oscillator) held on tempo by periodic recalibration against a hard physical reference — the fixed-cycle homing used in step-actuator positioners, where the mechanism is driven into a hard stop every N actuations to zero accumulated drift. Distributed over the cavity bus, this phases the array to a common beat. A graded ("chirped") guide, in which the local resonant frequency varies along the structure, can be written by a deposition-rate program during growth and retuned afterward by strain or applied field, supplying a position-to-frequency map for the channel layout.

Scope note: every primitive here builds storage, addressing, sensing, and signaling. None implies the device is "alive" or evolving; those claims are out of scope for this engineering paper and are treated separately (and far more skeptically) in the companion life paper.

Figures

Fig. 1 — Information-layer simulation. The over-binning cliff, the capacity-vs-dimension peak, and two-tier Arrhenius retention.




Fig. 1b — Closed-form capacity bound vs Monte Carlo (matches to ~ 0.001).




Fig. 1c — Net reliable bits vs dimension (peaks at four coupled axes; declines beyond).




Fig. 1d — Optimal radix collapse (12 $\rightarrow$ 3 $\rightarrow$ 2 as read-noise rises).




Fig. 1e — Coarse reading as a classifier (§4a): a lossy low-resolution read classifies category as well as or better than a full read at $\sim 60\times$ compression — the constructive twin of the over-binning cliff.




Fig. 2 — The signed alphabet (chirality sign $\times$ bag-charge magnitude).




Fig. 3 — Crossed acoustic/microwave addressing head.




Fig. A1 — Implementation primitives: memristor cell, self-reading inductor (emergent EMF), and triple-duty disclination transducer. (Schematic; full component grading in the Crystal EE Toolkit companion.)

References

All references verified against primary sources.

This framework (self-citations — the body of work this paper supports and extends; author-supplied DOIs, confirm resolution before posting): 1. Schulz, M. *Turtles, or Crystals* (Part One — one substrate in depth; the direct predecessor). Zenodo (2026). DOI: 10.5281/zenodo.20484438. 2. Schulz, M. *Turtles, All the Way Up* (the framework + its testable core). Zenodo (2026). DOI: 10.5281/zenodo.20486047. 3. Schulz, M. *Topological Linking-Number Heredity in a Crystalline Substrate, and the Open-Ended-Evolution Wall* (companion artificial-life paper, "Part Two"). This deposit (2026). 4. Schulz, M. *Comparative Biosonar & Navigation* (navigation-domain ancestor; relevant to addressing/homing). Zenodo (2026). DOI: 10.5281/zenodo.20502058.

External prior art: 3. Skyrme, T. H. R. A unified field theory of mesons and baryons. *Nucl. Phys.* (1962). 4. Mermin, N. D. The topological theory of defects in ordered media. *Rev. Mod. Phys.* 51, 591 (1979). 5. Burton, Cabrera & Frank. The growth of crystals and the equilibrium structure of their surfaces. *Phil. Trans. R. Soc. A* (1951). 6. Frank, F. C. The influence of dislocations on crystal growth. *Discuss. Faraday Soc.* (1949). 7. Foster, D., Kind, C., Ackerman, P. J., Tai, J.-S. B., Dennis, M. R. & Smalyukh, I. I. Two-dimensional skyrmion bags in liquid crystals and ferromagnets. *Nature Physics* 15(7), 655–659 (2019). DOI 10.1038/s41567-019-0476-x. 8. Zhang, Y. et al. Skyrmions based on optical anisotropy for topological encoding. *Light: Science & Applications* (2026). DOI 10.1038/s41377-026-02307-4 (arXiv:2508.16483). [The "60° rule": encoded information stays topologically protected provided the realized field differs from the design by up to 60° everywhere.] 9. Shechtman, D. et al. Metallic phase with long-range orientational order (quasicrystals). *Phys. Rev. Lett.* 53, 1951 (1984). 10. Kitaev, A. Fault-tolerant quantum computation by anyons (toric/surface codes). *Ann. Phys.* 303, 2 (2003). 11. Pappu, R., Recht, B., Taylor, J. & Gershenfeld, N. Physical One-Way Functions. *Science* 297(5589), 2026–2030 (2002); Gassend, B., Clarke, D., van Dijk, M. & Devadas, S. Silicon physical random functions. *Proc. 9th ACM CCS*, 148–160 (2002). 12. Auld, B. A. *Acoustic Fields and Waves in Solids* (Wiley, 1973; 2nd ed. Krieger, 1990) — standard reference for surface- and bulk-acoustic-wave devices. 13. Incumbent-memory benchmark figures (endurance, retention, cell area, bits/cell for NAND/DRAM/STT-MRAM/ReRAM) are drawn from standard device-technology reviews and comparison tables in the memory-engineering literature (e.g. surveys of emerging non-volatile memory; STT/SOT-MRAM endurance $\sim 10^{14}$ – 10^{16} and NAND $\sim 10^4$ cycles are widely reported). These are cited as representative industry figures for comparison, not as our measurements; specific values should be confirmed against current datasheets before external circulation.