

Acousto-Cymatic Phase-Change Memory: A Speculative Theoretical Framework for High-Density Data Storage in Confined Aqueous Media

Sahar Barak*

Independent Researcher

Preprint - June 11, 2026

Speculative “idea paper” (cf. physics.pop-ph/cs.ET).

Abstract

As conventional magnetic and solid-state storage technologies approach fundamental physical scaling limits - most notably the superparamagnetic limit [1, 2] - new architectures are sought to meet global data demand. This paper proposes, as a deliberately speculative thought experiment, a data-storage architecture we call *Acousto-Cymatic Phase-Change Memory* (ACPCM). By drawing engineering analogies from Heat-Assisted Magnetic Recording (HAMR) [3], three-dimensional charge-trap flash [5, 6], megahertz acoustic holography [12], and scientific data sonification [27], we sketch a method for encoding digital data into the physical topology of a confined, rapidly frozen aqueous “pod.” In the proposal, data is encoded as acoustic frequencies, organized into a three-dimensional grid by cymatic pressure nodes, and locked into a solid lattice by a HAMR-style “broadcast-and-freeze” phase transition; read-back is designed around established megahertz acoustic holography and neural-network field inversion, isolating the speculative burden on the write side. We are explicit throughout about which ingredients are established physics and which are contested or non-mainstream (“water memory,” structured-water chains, and W. Russell’s octave cosmology). Rather than asserting feasibility, we cast the architecture as a set of falsifiable predictions and identify the physical obstacles - above all the energetics, lattice distortion, and the ~ 50 fs memory loss and picosecond reorganization of the hydrogen-bond network [31] - that would have to be overcome for any such scheme to be viable. We make the reckoning quantitative: granting even the contested “structured-water” figures their own most favourable numbers, the binding constraint is not molecular packing but the *acoustic diffraction limit of the read path* (in ice, not water), which caps a realistic megahertz-addressed pod at $\sim 10^4$ – 10^5 bits per cubic millimetre - some 9–14 orders of magnitude below the molecular and structured-water ceilings - while the energy to freeze a bit lands orders of magnitude above incumbent storage [68]; against established archival media such as 5D fused-silica glass [40] and DNA [44], ACPCM is, on these numbers, uncompetitive. We further locate the proposal against the *sonocrystallization* literature [53]: ultrasound demonstrably imprints structure on freezing ice, but through *cavitation* - a stochastic, pattern-destroying mechanism - so the architecture’s cheapest test is at risk of a misleading positive and must control for it, while the one genuinely buildable variant (acoustically assembling seeded scatterers and locking them by solidification, already demonstrated [21]) records particle positions rather than the water field itself. Finally we sketch the pivot that lets the idea stand as a *device* rather than a high-density medium: built from directional (freeze-cast) solidification below the Mullins–Sekerka instability, seeded acoustic scatterers, and millitesla magnetic grain orientation - each demonstrated in isolation - and pitched not on density but on substrate cost, non-toxicity, and recyclability in environments where cold is free. The contribution is a conceptual synthesis, a quantitative bound, a device-shaped operating point, and a research agenda, *not* a validated device.

*Correspondence: hi@saharbarak.dev.

Keywords: unconventional computing; acoustic holography; acoustofluidics; phase-change media; ice nucleation; data sonification; speculative storage architectures.

This is a **speculative theoretical proposal**, written in the spirit of an “idea paper.” It combines *established* engineering mechanisms with *contested or refuted* premises, and it should be read as a hypothesis generator, not as a report of a working device or of accepted science.

- **Established** (well supported in the literature): the superparamagnetic limit and HAMR [1, 2, 3, 4]; 3D charge-trap flash [5, 6, 7]; chalcogenide phase-change memory [8, 9, 10]; passive acoustic-hologram phase plates and acoustofluidic patterning [12, 17, 20]; deep-learning acoustic-hologram synthesis [25]; and data sonification as a methodology [26, 27].
- **Contested / minority view**: “exclusion-zone / fourth-phase” structured water [35] (the interfacial effect is real but its structural interpretation is disputed [34]); quantum coherence in microtubules (Orch-OR) [32], challenged on decoherence grounds [33].
- **Refuted / pseudoscientific** (used here only as historical or motivational framing, *not* as evidence): “water memory” [36], which an on-site investigation found did not survive blinding [37] and which later revivals [38] did not reproduce; and W. Russell’s octave periodic system [39].

A claim’s appearance in this paper is *not* an endorsement of its validity. Crucially, the established citations ground only their own subsystems; none of them lends support to the central acousto-cymatic conjecture, which is ours and is clearly labelled as speculation throughout. Section 9 and table 2 make the dependencies explicit.

1 Introduction

The data-storage industry is constrained by a set of competing requirements - high areal density (which demands small grains), thermal (bit) stability (which demands high magnetic anisotropy), and writeability (which demands a switching field a head can actually supply) - that become mutually exclusive as bits shrink; in the later literature this is often called the *recording trilemma* [2]. The underlying physics is thermodynamic: as a grain’s volume V shrinks, its anisotropy energy $K_u V$ falls toward the thermal energy $k_B T$, and once the stability ratio $K_u V / k_B T$ drops below ~ 40 – 60 the magnetization reverses spontaneously - the *superparamagnetic limit*, whose onset Charap, Lu, and He first projected for thin-film media [1] and which Weller and Moser analyzed in its modern form [2].¹ The industry’s response, *Heat-Assisted Magnetic Recording* (HAMR), uses a near-field plasmonic transducer to heat a high-anisotropy medium close to its Curie temperature for well under a nanosecond, transiently collapsing its coercivity so a modest field can write a bit, after which rapid cooling “freezes” the magnetization into a thermally stable state [3]; the near-field-transducer write process has been demonstrated experimentally [4] and HAMR is now shipping commercially.

In parallel, solid-state drives escaped planar scaling by going vertical: 3D NAND with *charge-trap flash* (CTF) stores charge in a discrete, electrically insulating silicon-nitride layer rather than a continuous conductive floating gate, so a single local defect leaks only its neighborhood rather than draining a shared pool [5]; the move to stacked, “punch-and-plug” vertical strings decoupled density from lithography [6, 7].

This paper asks a deliberately unconventional question: *what if the two core ideas behind these technologies - rapid state-change writing (HAMR) and microscopic spatial isolation (CTF) - were applied to an exotic medium, namely water frozen on demand?* Liquid water is, on its face, the worst imaginable storage medium: its hydrogen-bond network loses memory of its configuration within ~ 50 fs and fully reorganizes

¹Their projected areal-density upper bound of ~ 36 Gbit/in² for longitudinal media was later surpassed by media engineering and perpendicular recording; we cite it as a projected estimate, not a proven wall.

on a picosecond timescale [31], so any imprint in the *liquid* is erased almost immediately. Our proposal confronts this directly by never trying to store information in the liquid: information exists only fleetingly as a broadcast acoustic field and is committed to memory by an *instantaneous phase change* into ice, in loose analogy to HAMR’s heat-then-quench write cycle [3] and to chalcogenide phase-change memory’s melt–quench amorphization [8, 9].

We stress at the outset (and repeat in section 9) that several premises in the draft framing of this idea - “water memory,” long-lived “structured-water” chains, and an octave-based re-mapping of the periodic table - are *not* part of accepted physics or chemistry; the first was specifically investigated and not reproduced [37, 36]. We include them transparently, label their status, and engineer the architecture so that the refuted ideas carry no logical load. The intended contribution is a clean conceptual synthesis across fields that are rarely discussed together, plus an explicit map of where the idea would break.

Where this sits among unconventional storage. The search for post-magnetic archival media is well populated, and ACPCM should be judged against it. *5D optical data storage* writes birefringent nanogratings into fused silica with a femtosecond laser - three spatial coordinates plus slow-axis orientation and retardance - with a peer-reviewed room-temperature decay time of order 10^{20} years and thermal stability to 1000 °C [40], and demonstrated high-speed writing toward hundreds-of-terabytes-per-disc capacities [41]. *DNA data storage* encodes bits in synthesized oligonucleotides: from the first megabit-scale demonstrations [42, 43] to a coding density of 1.57 bits/nucleotide and an experimentally realized 215 PB/g [44], with random access at the ~ 200 MB scale [45] and a molecular ceiling near 10^{18} bytes/mm³ [46]. *Volume holographic storage* records data as superimposed interference gratings with a theoretical ceiling of order V/λ^3 [47, 48]. Each is mature or at least experimentally grounded; ACPCM is far more speculative than any of them, and section 6 places it on the same axes so the comparison is explicit rather than rhetorical.

Contributions.

1. A subsystem-by-subsystem architecture (medium, format, write, read) that transposes HAMR and CTF design principles onto a rapidly frozen aqueous medium (sections 2 to 5).
2. An explicit *epistemic ledger* (table 2) separating established mechanisms from speculative ones, so the proposal cannot be mistaken for a validated device and so that mainstream citations are not laundered into supporting fringe claims.
3. A *quantitative reckoning* (section 6) that derives the achievable density from the acoustic diffraction limit, takes the contested structured-water capacity figures at face value to show what would be required, and benchmarks the resulting density and energy-per-bit against 5D glass, DNA, and incumbent semiconductor storage (table 1).
4. A set of falsifiable predictions and order-of-magnitude obstacles (sections 8 and 9) that define what an experimental test would have to measure.

2 The storage medium: confined aqueous “pods”

For water to act even momentarily as an addressable medium, it cannot be a bulk volume. In bulk liquid, distinct driving frequencies superpose and produce beat phenomena, so independently addressed “cells” would modulate and erase one another (crosstalk). The proposal therefore confines the medium to discrete *pods* that physically isolate addressable regions - an acoustic analogue of CTF’s discrete charge traps, where isolation, not bulk capacity, is the design goal [5].

A note on scale is needed up front, because it disciplines everything that follows. The addressing grid is drawn by a megahertz field, so the smallest feature it can write or read is set by that field’s wavelength: at 5 MHz to 50 MHz the half-wavelength in water is ~ 15 μm to 150 μm (section 6.2). A pod must therefore be

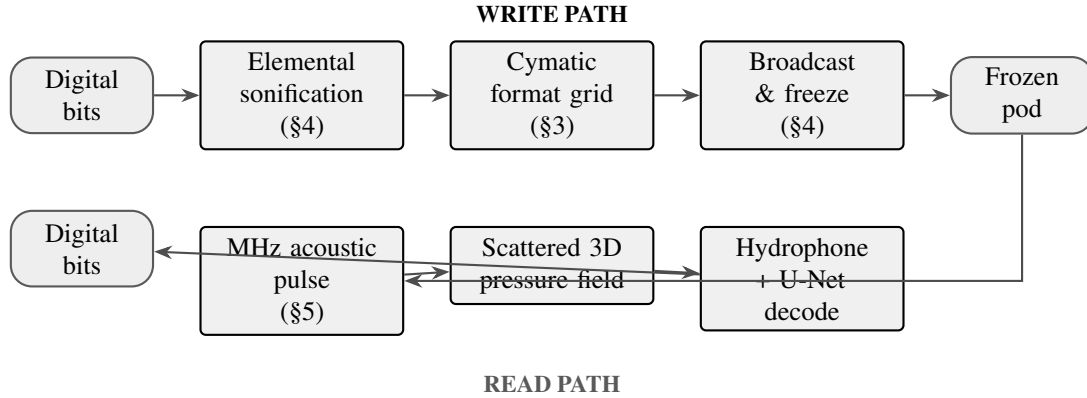


Figure 1. End-to-end ACPCM data path. The write path encodes bits as acoustic tones, imposes a cymatic addressing grid, and commits the field to a solid lattice by freezing. The read path probes the frozen pod with a megahertz pulse, treats the ice topology as an acoustic phase plate, and decodes the scattered field with a neural network. The read path and the write *actuation* reuse established techniques; the speculative weight is isolated in the freeze step (Assumption 4). Each block’s epistemic status is given in table 2.

at least a few wavelengths across to host an internal grid at all - of order 0.1 mm to 1 mm, not sub-micron. We retain “microscopic” only in the loose sense of sub-millimetre; the isolation in Assumption 1 is between neighbouring pods in an addressed array, not within a single nanoscale cell. Any nanometre-scale structure invoked below (microtubule water, structured-water “pearls”) is therefore *motivational only* and explicitly *not* the addressed unit - a point the capacity reckoning of section 6 makes quantitative.

Assumption 1 (Confinement enables isolation). Sub-millimetre pod confinement suppresses inter-pod acoustic crosstalk enough that independently addressed pods do not annihilate one another during the write window.

As biological *motivation only*, we note the recurring (and contested) claim that cytoskeletal microtubules - hollow cylinders with an inner diameter of order 15 nm - might shield ordered internal water from thermal noise, a premise of the Orch-OR quantum-biology hypothesis [32]. (This is a nanometre-scale intuition pump, four orders of magnitude below the sub-millimetre addressed pod above; we borrow only the *qualitative* notion that confinement protects order.) That premise is itself disputed: Tegmark’s decoherence estimates put any quantum coherence in warm, wet microtubules many orders of magnitude below cognitive timescales [33] (a critique that proponents have in turn disputed on modelling grounds). We therefore use microtubules strictly as an intuition pump for “confinement protects order” and assert nothing about biological memory. Likewise, the notion of long-lived spherical “water pearls” or “trimmed chains” as information carriers belongs to the structured-water literature [35]; the underlying interfacial exclusion-zone effect is real and replicated, but its “fourth-phase” structural interpretation is a contested minority view that conventional mechanisms may explain better [34]. In ACPCM such structures, if they exist at all, are relevant only *after* freezing has removed the liquid’s thermal lability, not before.

3 Format: a cymatic cylinder–head–sector grid

Conventional drives impose logical structure - tracks and sectors mapped to logical block addresses - onto a physical medium. ACPCM proposes to impose this structure *acoustically*. Using megahertz ultrasound projected into the pod, intersecting wavefronts establish a three-dimensional pattern of standing-wave nodes (pressure minima) and antinodes. The nodal surfaces are intended to act as “sector walls,” partitioning the pod into addressable voxels with definite spatial coordinates.

The *patterning of matter by sound* is established physics. Since Faraday it has been known that vibrating surfaces and liquid layers drive particles into ordered figures [11]; modern acoustofluidics uses MHz standing waves to focus, separate, and trap particles and cells at radiation-force nodes [20], down to one cell per well [19], and single-sided holographic arrays form programmable “acoustic tweezers” and vortex traps that levitate and manipulate discrete objects [17, 18]. Passive acoustic-hologram phase plates can sculpt remarkably intricate 3D pressure landscapes from a single transducer [12, 13].

The *spatial arrangement of discrete particles, droplets, or cells at acoustic nodes* is established and reproducible (fig. 2) - but it is transient (it lasts only while the field is on) and it organizes *pre-existing objects*. The stronger claim that such fields coerce *water molecules themselves* into a persistent “supramolecular network” - sometimes labelled “molecular cymatics” - is **not** established: it has no accepted mechanism and no primary literature specifically linking acoustic-field forcing to supramolecular water ordering (the structured-water literature [35, 34] concerns interfacial proximity, not acoustic induction), and liquid-water hydrogen bonds relax in picoseconds [31]. It is treated here purely as a hypothesis (Assumption 2), and must not be confused with the established assembly result.

Assumption 2 (Cymatic addressing). A megahertz standing-wave field can define stable, reproducible nodal boundaries at the length scale of the intended voxels, and these boundaries survive the freezing transition well enough to serve as a read-time coordinate system.

3.1 Two-axis addressing: slice into discs, then encode each disc

We separate *addressing* from *encoding* by superposing two orthogonal acoustic fields of very different scale - directly analogous to a hard drive’s split between the coarse mechanical track/cylinder geometry and the fine, high-frequency bit cells written along a track.

1. **A low-frequency axial “slicing” wave (the disc former).** A single longitudinal standing wave along the pod’s z -axis, at the *low* end of the band, partitions the pod into a stack of parallel nodal planes spaced by $\lambda_z/2$. Its pressure antinodes define a register of flat “discs” (the platter/cylinder analogue) - a small number of robust, widely spaced, easily resolved layers. Because λ_z is large, these boundaries are the most likely structure to *survive freezing* (section 9) and so anchor the read-time coordinate system.
2. **An orthogonal directional “encoding” wave (the in-disc data).** Within each disc, a transverse, higher-frequency, *travelling-or-standing* field projected along (x, y) writes the payload: the fine pressure pattern in the plane of a disc is the data, the multi-tone codebook of section 4 modulating amplitude and phase across the disc. The encoding wave rides at the fine end of the resolution budget; the slicing wave guarantees that the decoder always knows *which* disc a given fine feature belongs to.

The two axes are deliberately decoupled in frequency so that the coarse layer index is robust even where the fine payload is degraded - a graceful-degradation property borrowed from how a drive can still seek a track whose individual bits are marginal. It also makes the diffraction budget legible: the disc count is set by λ_z , the in-disc bit count by the transverse wavelength, and the achievable product is bounded in section 6.

Assumption 3 (Separable addressing). A coarse axial standing wave and a fine transverse field can be imposed simultaneously without the strong (low-frequency) former corrupting the weak (high-frequency) payload, and the disc boundaries remain registered to the in-disc pattern through the freeze.

4 Encoding and writing: the “broadcast-and-freeze” mechanism

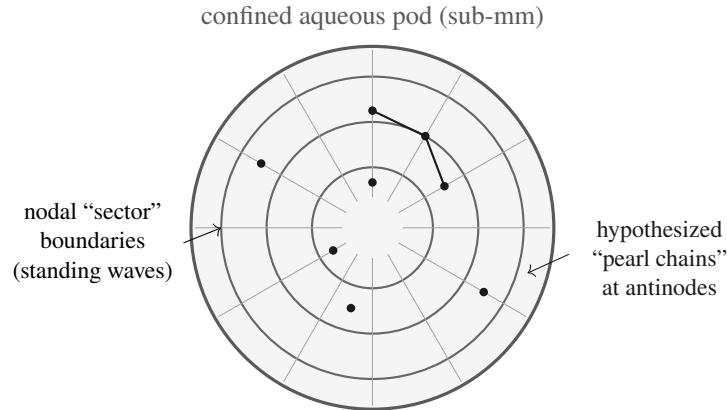


Figure 2. Schematic of a pod with a cymatic cylinder–head–sector (CHS) analogue. Concentric and radial standing-wave nodes (blue) partition the pod into addressable voxels - the established, but transient, particle-patterning physics. The hypothesized “pearl chains” resident at antinodes (amber) belong to the contested structured-water premise (Assumption 2). Not to scale.

4.1 Encoding by elemental sonification

Digital payloads are mapped to acoustic spectra. *Data sonification* - the principled mapping of data to non-speech sound - is a mature methodology with an NSF-commissioned status report [26], a canonical handbook [27], and a formal taxonomy whose defining criterion is that the sound reflect *objective* properties of the data through a systematic, reproducible transform [28]. Mapping the wavelengths of atomic emission lines into the audio band is a documented technique used for analysis, accessibility, and education - e.g. the audification of astronomical spectra [29] and an interactive sonified periodic table [30]. ACPCM uses such a mapping as a codebook: payload symbols index a dictionary of multi-tone acoustic “words.” These audio-band tones serve as *codebook symbols*, not literal drive frequencies: the transducer operates in the MHz band (section 3), and each symbol indexes a distinct multi-tone pattern within it.

It is essential to be precise about what sonification *is not*. The light-to-audio mapping (from $\sim 10^{14}$ Hz to $\sim 10^2$ – 10^4 Hz) is an arbitrary, reproducible rescaling chosen for human perceptibility [30]; it is *not* a physical resonance, it reveals no “hidden physics,” and it carries no information beyond what is already in the source data [28]. The draft framing additionally invokes W. Russell’s octave cosmology, in which the elements are arranged on a spiral of ten octaves balancing “centripetal” and “centrifugal” tendencies [39]. We are explicit that this is a **historical/esoteric** system - metaphysics, not chemistry - used here, if at all, only as one arbitrary-but-fixed codebook among many. Nothing in the architecture depends on Russell’s metaphysics being correct; any injective, noise-separated frequency code would serve, and a *designed* code (e.g. orthogonal frequency-division multiplexing) is strictly preferable.

4.2 Writing by phase change

The pod is held near its freezing point. The transducer simultaneously broadcasts (i) the cymatic formatting field of section 3 and (ii) the data tones. While the field is held, the pod undergoes an *instantaneous phase transition* to a solid. In analogy to the melt–quench programming of chalcogenide phase-change memory [9, 10] - and, more loosely, to HAMR’s heat-and-quench write cycle [3] - in which a pulse melts a bit and rapid cooling freezes it into a distinct solid phase, the rapid solidification is hypothesized to “record” the instantaneous acoustic field as a frozen-in topology - displacements, density modulations, and interface positions locked into the lattice.

Assumption 4 (Faithful commit). Freezing transfers the broadcast acoustic field into a static, readable lattice topology with sufficient fidelity and reproducibility to be decoded, and the transition’s own dynamics (stochastic nucleation, latent-heat release, ~9% volumetric expansion, dendritic growth) do not destroy the encoded pattern.

Assumption 4 is the architecture’s load-bearing physical claim and its most likely point of failure; we return to it in section 9. We emphasize that the HAMR and PCM citations establish only that *their own* rapid-quench mechanisms work; they provide *no* evidence that an acoustic field can be recorded by freezing water. That bridge is ours, and it is conjecture.

4.3 What the freezing-under-sound literature actually says

There *is* an established body of work in which ultrasound applied during freezing leaves a structural imprint in the resulting ice - *sonocrystallization*, or ultrasound-assisted freezing (UAF), studied extensively in food and pharmaceutical processing [53, ?, 54]. Power ultrasound demonstrably triggers ice nucleation in supercooled water, shortens freezing time, and controls ice-crystal size and dendrite morphology. At first glance this is encouragement for Assumption 4: sound applied during a freeze *does* change the solid that results.

The mechanism, however, cuts the wrong way, and honesty requires stating it plainly. The established coupling is *cavitation*: the field nucleates and collapses gas/vapour bubbles, and those bubbles - not the field’s spatial pattern - seed the ice [53, 55]; collapsing bubbles further emit microjets and shock waves that *fragment* growing dendrites into secondary nuclei [53]. This is spatially *stochastic* (bubble sites are random) and *pattern-destroying* (fragmentation erases fine structure), the opposite of a faithful field-to-lattice transcription. So the one real sound→ice channel supports the weak claim (a patterned freeze differs from a sham freeze) while actively undermining the strong one ACPCM needs (the difference *is* a readable copy of the field). Section 9 draws the consequence: the cheapest test, Prediction 1, is at risk of a *misleading positive* driven by cavitation rather than recording, and must be designed to exclude it.

4.4 A buildable special case: acoustically assembled, then locked

One variant of the write step is *not* speculative at all, and we flag it as the project’s honest engineering exit. If the pod is seeded with discrete, acoustically responsive tracer objects (micro-particles, beads, or cells), then arranging them at the standing-wave nodes and *locking* them in place by solidifying the surrounding medium is demonstrated, current technology: holographic ultrasound fields assemble particles, gel beads, and cells into prescribed 3D shapes that are then fixed permanently by gelation of the host medium [21], and acoustic focusing of beads/cells in a droplet followed by cross-linking locks their positions to within ~10% of the target line [22]. Substitute “freeze” for “gel” and this is precisely ACPCM’s write–lock cycle, minus the speculative leap.

The price of buildability is that this records the *positions of seeded objects*, not the topology of the water itself - it is a particle-in-matrix memory, not “molecular cymatics.” It still inherits the acoustic read limit of section 6.2 (the locked pattern is read no finer than the ice wavelength), so it does not rescue the density verdict; and it is a more modest claim than the pure-water proposal. We include it because it draws the boundary exactly: everything up to and including *assemble-and-lock discrete scatterers* is established; the single remaining conjecture is that the same freeze records the *continuous water field* with no seeded objects at all (Assumption 4).

5 Reading: megahertz acoustic holography

Once committed to ice, the data lives in microscopic features of the solid, far below the resolution of any microphone. Read-back is posed as an *acoustic holography* problem [12].




Figure 3. Read chain. A megahertz pulse scatters off the frozen pod, whose topology acts as an acoustic phase plate [12]; the resulting volumetric field (scalar pressure plus, optionally, particle-velocity components) is sampled by a hydrophone array [23, 24] and inverted to a bitstream by a trained network [25]. This read chain reuses only established techniques; the speculative burden lies entirely on the write side (Assumption 4).

The ice as a phase plate. A megahertz ultrasonic pulse illuminates the frozen pod. (That an acoustic probe can interrogate the water-to-ice transition at all is established: MHz plate waves resolve the freezing transition through its attenuation signature [56].) Because acoustic holography is diffraction-limited, resolving microscopic features requires correspondingly short wavelengths - i.e. MHz (or higher) ultrasound, with millimetre-to-sub-millimetre wavelengths in water and ice [12, 14]. The ice topology then acts as a passive *acoustic phase plate* [12, 15]: the scattered wave is transformed into a complex, three-dimensional volumetric pressure field that encodes the structure it bounced off, and such plates can in principle be stacked or reconfigured to multiplex fields [16].

Capturing the field. The scattered field is sampled by a calibrated needle hydrophone over a measurement surface, the standardized method for spatially mapping ultrasound pressure [23]. We note an honest caveat from the metrology literature: while focal-pressure measurements can be highly *repeatable* (coefficient of variation below 2%), inter-hydrophone *agreement* for tightly focused MHz fields can differ by tens of percent due to spatial averaging and frequency-response uncertainty [24] - a real ceiling on read fidelity. A frequently proposed enhancement is *vectorial* sensing: recording not only the scalar pressure p but the particle-velocity components (v_x, v_y) . We flag, however, that in a fluid these are *not* independent degrees of freedom - they are fixed by spatial gradients of the same scalar pressure via Euler’s relation - so they add an information *dimension* under a physical constraint rather than free, independent bandwidth; this remains an emerging idea, not a settled gain.

Neural decoding. Finally, the measured field is inverted to the original bitstream by a deep network - e.g. a U-Net of the kind already used for fast acoustic-hologram synthesis [25]. The network *architecture* is mature, but the *inverse* field-to-source decoding ACPCM requires is harder and less established than the forward synthesis those works demonstrate: such inverse problems are ill-posed and a network can hallucinate plausible outputs, so any demonstration must include held-out, information-theoretic controls (Prediction 3).

6 Capacity, density, and energy: a quantitative reckoning

An architecture proposal that omits the numbers is untestable as engineering. Here we put ACPCM on the same axes as real storage. The conclusion is unflattering and, we think, clarifying: the binding limit is not how finely water can be arranged but the wavelength of the wave used to read it, and on that limit ACPCM is many orders of magnitude short of every incumbent.

6.1 The optimistic ceilings (granted, then discounted)

A 1 μL pod (a 1 mm cube of water, ~ 1 mg) holds $\sim 3.3 \times 10^{19}$ molecules. A naive “one bit per molecule” bound therefore suggests $\sim 3 \times 10^{19}$ bits (≈ 4 EB) *per microlitre* - a physically meaningless upper bound, but a useful ceiling.




Figure 4. Volumetric storage density across media (log scale; figures as in table 1). The read-limited ACPCM reality (bottom, this work) sits roughly nine orders of magnitude below the structured-water theory’s own claim and fourteen below the one-bit-per-molecule ceiling, and below every established medium. The binding constraint is the ice-borne read wavelength (section 6.2), not how finely matter can be arranged.

The structured-water literature supplies a more specific, internally quantified claim, which we state explicitly and *without endorsement* (it rests on the contested “fourth-phase” premise [35, 34]). In that account the storage unit is not a molecule but a ~ 15 nm “pearl” of $\sim 85,000$ molecules; pearls link into ~ 300 nm “chains” of ~ 20 pearls, with chains below ~ 100 nm deemed too short to resolve a distinct signature. Treating one chain (a ~ 15 nm-wide, ~ 300 nm-long cylinder, $\sim 5.3 \times 10^4$ nm³) as one symbol gives $\sim 1.9 \times 10^{16}$ chains/cm³, i.e. the literature’s own headline of $\sim 1.9 \times 10^{16}$ bytes/cm³ (≈ 19 TB per microlitre, $\sim 700\times$ the density of 5D glass) [35]. We carry this figure forward only to show, next, that even granting it in full changes nothing: a constraint downstream of storage dominates.

6.2 The diffraction limit sets the real ceiling

Read-out is acoustic (section 5) and - crucially - propagates through the frozen *ice*, where the sound speed is $c \approx 3200$ m s⁻¹, roughly twice the 1500 m s⁻¹ of liquid water. Spatial resolution is bounded by the probe wavelength *in that medium*: a 50 MHz read pulse has $\lambda \approx 64$ μ m in ice and resolves voxels no finer than $\lambda/2 \approx 32$ μ m. (Features may be *written* more finely - the cymatic field forms in the liquid, where $\lambda/2 \approx 15$ μ m - but they cannot be *read* below the ice limit, so the channel is read-limited.) Absorption climbs as $\sim f^2$, so halving the voxel by going to ~ 100 MHz costs steeply in penetration depth - a hard, physics-imposed resolution–penetration trade. A 1 mm pod thus contains $\sim (1000/32)^3 \approx 3 \times 10^4$ resolvable voxels; allotting several bits per voxel by exploiting phase and amplitude lifts this only to $\sim 10^5$ bits - of order 10 kB per cubic-millimetre pod.²

The verdict follows immediately. The realistic $\sim 10^5$ bits/mm³ sits roughly *nine* orders of magnitude below the structured-water theory’s own $\sim 1.5 \times 10^{14}$ bits/mm³, and roughly *fourteen* below the one-bit-per-molecule ceiling. *You cannot read what you cannot resolve*; the molecular real estate is irrelevant if the only available probe paints it with a ~ 32 μ m brush. This is the quantitative core of why ACPCM, even with every contested premise granted, is not a high-density medium.

²The two-axis scheme of section 3.1 makes the budget concrete and, deliberately, more conservative: a 5 MHz axial former (liquid-phase $\lambda_z \approx 300$ μ m) writes ~ 6 coarse, freeze-robust discs across a 1 mm pod - well above the ice read limit, so all are resolvable - while the read-limited transverse cells number $\sim (1000/32)^2 \approx 10^3$ per disc, for $\sim 6 \times 10^3$ voxels: below the isotropic $\sim 3 \times 10^4$ ceiling, the price of trading resolution for a coarse, freeze-robust layer index.




Figure 5. Write energy per bit (log scale). Freezing a $\sim 1 \mu\text{L}$ pod and amortizing over its $\sim 10^5$ read-limited bits costs $\sim 4 \mu\text{J}$ of heat removed per bit (section 6) - five to seven orders of magnitude above incumbent picojoule-scale operations [68], and that is before refrigeration overhead, transducer drive, and neural decoding.

Medium	Volumetric density (order)	Status / note
5D fused-silica glass	$\sim 26 \text{ TB/cm}^3$	Emerging; $\sim 10^{20}$ -yr decay, 1000°C stable [40, 41]
DNA (synthesized oligos)	$\sim 215 \text{ PB/g}$ ($\sim 10^{17} \text{ B/cm}^3$); ceiling $\sim 10^{18} \text{ B/mm}^3$	Emerging; random access demonstrated [44, 46]
Volume holography	$\sim 1/\lambda^3$ (order 1 TB/cm^3)	Niche; theoretical bound [47, 48]
ACPCM - structured-water accounting	$\sim 1.9 \times 10^{16} \text{ B/cm}^3$ ($19 \text{ TB}/\mu\text{L}$)	Speculative; contested premise [35]
ACPCM - diffraction-limited reality	$\sim 10^7 \text{ B/cm}^3$ ($\sim 10^5 \text{ bits/mm}^3$)	This work; set by ice read wavelength (section 6.2)

Table 1. ACPCM against unconventional storage media on volumetric density. DNA per- cm^3 assumes a dry density $\sim 1.1 \text{ g cm}^{-3}$. The two ACPCM rows bracket the gap between what the contested premise claims and what the acoustic read path can actually resolve: ~ 9 orders of magnitude, with a further ~ 5 to the one-bit-per-molecule ceiling. The realistic ACPCM density is below every established medium in the table.

6.3 Energy and throughput

Committing a bit means freezing a droplet. Removing heat from a pod of mass m held ΔT above freezing costs at least $E = m(c_p \Delta T + L_f)$, with $c_p \approx 4.18 \text{ J g}^{-1} \text{ K}^{-1}$ [50] and $L_f \approx 333.5 \text{ J g}^{-1}$ [49]. For a $1 \mu\text{L}$ pod cooled 20 K then frozen, $E \approx 1 \times 10^{-3} \text{ g} \times (84 + 333.5) \text{ J g}^{-1} \approx 0.42 \text{ J}$. Spread over the diffraction-limited $\sim 10^5$ bits, that is $\sim 4 \times 10^{-6} \text{ J per bit}$ ($\sim 4 \mu\text{J/bit}$) of heat that must be *removed* - by a refrigerator with coefficient of performance well below unity, multiplying the wall-plug cost several-fold, and before counting transducer drive, hydrophone scan, and neural inversion. Incumbent operations sit at the picojoule-per-bit scale (on-chip and DRAM access; 3D-NAND programming amortizes to $\sim \text{pJ}$ –tens-of-pJ per bit) [68]. ACPCM's write energy is therefore $\sim 10^5$ – 10^7 \times worse per bit before refrigeration overhead. Throughput compounds it: freezing is a stochastic-nucleation event on millisecond-and-slower scales against HAMR's sub-nanosecond per-bit write [3], and because rewriting demands re-melting, the medium is closer to write-once archival than to random-access memory.

7 Epistemic ledger

Table 2 is the paper's central honesty device: each subsystem is paired with the real-world mechanism it borrows and with a blunt status label. The pattern is deliberate: *the read path and the write actuation are built from established techniques, while the speculative weight is concentrated in a small number of clearly marked assumptions about how freezing water records and preserves an acoustic field*. The refuted ingredient

ACPCM subsystem	Borrowed real mechanism	Status of the leap
Rapid “write” by phase change (section 4)	HAMR heat–quench; chalcogenide PCM melt–quench switching [3, 9]	Speculative (Assumption 4)
Assemble-and-lock seeded scatterers (section 4.4)	Holographic acoustic assembly fixed by gelation; acoustic focusing + cross-link [21, 22]	Established for <i>particles</i> ; the leap is pure-water recording
Sound→ice structural imprint (section 4.3)	Sonocrystallization / ultrasound-assisted freezing [53, 54]	Imprint established but via cavitation (stochastic, pattern-destroying)
Ordered ice by directional freeze (section 10)	Freeze-casting; planar-front stability; magnetic grain alignment [64, 63, 65]	Established at scale; the device regime, not the chaotic quench
Cell isolation via pods (section 2)	Charge-trap flash discrete traps [5]	Plausible (Assumption 1)
Cymatic addressing grid (section 3)	Acoustofluidic patterning; holographic tweezers [20, 17]	Particle nodes established ; molecular ordering speculative (Assumption 2)
Two-axis disc addressing (section 3.1)	Drive track/sector split; axial vs. transverse standing waves	Geometry plausible ; survival through freeze speculative (Assumption 3)
Sonification codebook (section 4)	Data-sonification methodology [26, 27]	Mapping established ; Russell octaves esoteric/optional [39]
MHz acoustic-holography read-out (section 5)	Acoustic holograms; hydrophone metrology; DL inversion [12, 24, 25]	Established techniques
“Structured-water” carriers (section 2)	EZ/fourth-phase water; Orch-OR water [35, 32]	Contested / minority [34, 33]
“Water memory” (motivation only)	Benveniste high-dilution claims [36]	Refuted [37]

Table 2. Epistemic ledger mapping each ACPCM subsystem to its borrowed real mechanism and the status of the conceptual leap. The architecture is engineered so that the *refuted* ingredient (water memory) carries *no* load - it is motivational framing only - while the genuinely novel risk is isolated in Assumption 4. Established citations support only their own row and lend no credibility to the speculative rows.

(“water memory”) is load-bearing nowhere; it is motivational framing only.

8 Falsifiable predictions

A speculative proposal earns its keep by being testable. ACPCM implies concrete, disprovable predictions; any one of them failing falsifies the corresponding assumption.

Prediction 1 (Field-to-lattice transfer). Freezing a thin water film while broadcasting a fixed MHz pressure pattern produces a *reproducible*, pattern-dependent modulation in the resulting ice (e.g. in density, birefringence, or scattered-field signature) that is absent under sham (unmodulated) freezing. *Test*: interferometric or

acoustic readout of patterned vs. control freezes under blinding; *falsified if* no above-noise, pattern-correlated structure survives (Assumption 4).

Mandatory confound control. A bare positive here is **not** sufficient, because sonocrystallization guarantees that *any* insonated freeze differs from a quiescent one through cavitation-seeded nucleation and dendrite fragmentation (section 4.3) [53] - an effect that carries no field information. The test must therefore (i) suppress cavitation as a confound (degassed water, elevated static pressure, or amplitudes below the cavitation threshold), and (ii) demand that the imprint *tracks the imposed pattern's geometry* - changing with node spacing and orientation and registering to the antinode lattice - rather than merely exceeding the sham. A difference that does not move with the pattern is a cavitation artefact, not a recording, and counts as a *falsification*, not a success.

Prediction 2 (Addressable nodes). The post-freeze structure is spatially organized at the wavelength scale of the imposed standing wave. *Falsified if* the frozen modulation shows no correlation with the node geometry (Assumption 2).

Prediction 3 (Decodable channel). A network trained on (pattern \rightarrow scattered-field) pairs recovers the input code above chance on held-out patterns, with a measurable channel capacity. *Falsified if* decoding accuracy is at chance on held-out data, indicating no recoverable information (and guarding against the network merely memorizing or hallucinating).

Prediction 4 (Diffraction-bounded, two-axis capacity). The recoverable channel capacity per pod scales as $\sim(D/\lambda)^3$ in the pod diameter D and read wavelength λ , and a low-frequency axial former produces a separable layer index (Assumption 3): the decoder recovers *which disc* a feature occupies independently of the in-disc payload. *Falsified if* the layer index is not separable, *or* - the genuinely exciting failure - if recovered capacity *exceeds* the acoustic diffraction ceiling of section 6.2, which would imply a sub-wavelength read channel and overturn the pessimistic density verdict.

If Prediction 1 fails - the most likely outcome given current physics - the architecture collapses to its established components, and that negative result is itself worth reporting.

9 Discussion: obstacles, limits, and why this is hard

We do not claim ACPCM is feasible. The honest reading is that it is *probably not*, for reasons worth stating precisely because they sharpen the predictions above.

The freezing transition is violent. On freezing, water releases its latent heat of fusion, $L_f \approx 333.5 \text{ J g}^{-1}$ [49], and expands by $\sim 9\%$ as it drops from $\sim 1.0 \text{ g cm}^{-3}$ to the $\sim 0.917 \text{ g cm}^{-3}$ of hexagonal ice I_h [50]; nucleation is stochastic and growth is dendritic. Any acoustic “imprint” must survive this disordering, mechanically disruptive event. A $\sim 9\%$ volume change (a $\sim 3\%$ linear strain) at a voxel scale of $\lambda/2 \approx 15\text{--}75 \mu\text{m}$ (10–50 MHz in water) implies linear displacements of order $\sim 0.5\text{--}2 \mu\text{m}$ - a non-negligible fraction of the intended nodal spacing itself, directly threatening the registration the two-axis scheme depends on (Assumptions 3 and 4).

The obvious escape - vitrify rather than crystallize - is itself punishing. Forming amorphous (glassy) ice instead of dendritic I_h avoids the worst of the crystallization disruption, but pure water vitrifies only above a critical cooling rate recently measured at $\sim 6.4 \times 10^6 \text{ K s}^{-1}$ [52] (first demonstrated by hyperquenching micron-scale droplets [51]), practical only for very small volumes - in direct tension with the $\sim 15 \mu\text{m}$ voxels the read path can resolve (section 6.2). Crystallize instead and you race a moving front: dendritic ice tip velocities reach $\sim 0.1\text{--}10 \text{ cm s}^{-1}$ at moderate supercooling [66, 67], fast enough to sweep a $\sim 15 \mu\text{m}$ voxel in microseconds and rearrange whatever the field had arranged. Homogeneous nucleation does not even begin until the classical limit of $\sim 235 \text{ K}$ ($\sim -38^\circ\text{C}$), so the “hold the field, then freeze” write cycle must thread

a narrow window of deep supercooling. Assumption 4 is thus in direct tension with the thermodynamics of solidification on every axis - rate, habit, and strain - and controlling nucleation site, rate, and crystal habit (I_h vs. vitrified) is a hard, partly unsolved problem in its own right.

The only established sound→ice channel is the wrong one. It would be convenient if the sonocrystallization literature (section 4.3) [53, 54] supplied the missing mechanism. It does not. That literature couples sound to ice through *cavitation* - bubble nucleation and collapse - which seeds ice at *stochastic* sites and fragments dendrites [53, 55]; it changes ice statistics, not ice *coordinates*. A faithful commit (Assumption 4) needs the inverse: a deterministic, spatially registered transcription of a continuous field. The established effect is thus simultaneously the best evidence that freezing-under-sound does *something* and the clearest reason to expect that something to be informationally empty - which is exactly why Prediction 1 carries a mandatory cavitation control.

What the nearest experiments actually show. No one has performed the exact ACPCM experiment - impose a cymatic standing wave in bulk water, freeze, and recover spatial information - but the closest containerless-freezing work is consistent on the point that matters. When a water droplet is frozen while held in an acoustic levitation field, it supercools, nucleates a crust, and then solidifies by *inward dendritic growth*; the field's role is to *hold* the droplet without a container, not to write its interior, and the resulting ice records the thermal field and crystallography, not the acoustic node geometry [57]. Where an imposed field *does* leave a preserved pattern, it does so by arranging *discrete objects*: suspended particles in a levitated or drying droplet are driven to the pressure field's nodes and that arrangement survives into the dried/solid granule [58]. Both observations reproduce the boundary of section 4.4: fields move and freeze *things*; ordinary crystallization writes the *ice*.

Field-stabilised (containerless) freezing: a better rig, not a better device. A natural instinct is to hold each pod still and freeze it in place by a static field rather than against a wall. This is real and worth using. Water is diamagnetic and can be *levitated* in a strong gradient magnetic field - the route by which Tagami and co-workers performed containerless solidification of magnetically levitated water in a hybrid magnet [59], and which a recent design achieves for ~0.3 mm droplets with only ordinary permanent magnets [60]; charged droplets are likewise held contact-free in electrostatic levitators, the standard tool for deep-undercooling solidification studies [61]. Static and oscillating magnetic fields also demonstrably bias ice nucleation and crystal size [62], and the magnetic anisotropy of ice even lets a strong field orient the crystal *c*-axis - a genuine lever on habit and registration. The honest verdict is that this improves the *experiment* rather than the architecture: containerless levitation removes wall (heterogeneous) nucleation, enables the deep, controlled supercooling Assumption 4 needs, and is in fact the cleanest rig in which to run Prediction 1. But it does not address the load-bearing failure - the freeze front still overruns the pattern (above), the read path is still diffraction-limited (section 6.2), and a tesla-scale magnet or kilovolt trap *per pod* runs directly against the cheap-archival premise that motivates the medium. Levitation makes the test honest; it does not make the device feasible.

Liquid water has no long-term memory. The hydrogen-bond network reorganizes on picosecond timescales [31]; there is no credible evidence for persistent information storage in *liquid* water, and the specific "water memory" claims invoked as motivation were investigated and not reproduced [37, 36], including later revivals [38]. ACPCM only sidesteps this objection by storing nothing in the liquid - which places the entire burden on the freeze step.

Resolution versus penetration. Resolving microscopic features acoustically requires very high frequencies, where absorption in water and ice rises steeply; there is a real, physics-imposed tension between resolution and penetration that caps achievable density. Read-out is also an ill-posed inverse problem, and hydrophone metrology already shows tens-of-percent inter-device disagreement for tightly focused MHz fields [24], so any demonstration must include rigorous, blinded, information-theoretic controls (Prediction 3).

Thermodynamic and energy cost. Writing a bit by freezing and reading it by illumination implies an energy budget per bit that is plausibly orders of magnitude worse than incumbent technologies, before considering throughput. Even if every assumption held, the density/energy/latency envelope would need to beat HAMR [3] and 3D NAND [7] to matter - a very high bar.

What would change our mind. A positive, replicated result on Prediction 1 under blinded controls *with cavitation excluded* - pattern-dependent, above-noise, geometry-tracking structure that survives freezing and carries recoverable information - would promote Assumption 4 from speculation to phenomenon and justify pursuing the rest. Absent that, the honest fallback is not nothing: the particle-seeded, assemble-and-lock variant of section 4.4 is already demonstrated in adjacent fields [21, 22] and is the buildable core that survives if the pure-water conjecture fails. What dies in that case is specifically the claim that water records a continuous acoustic field without seeded scatterers; what remains is a modest, read-limited particle-in-ice memory and an organizing metaphor for cross-disciplinary analogy, not a competitive storage technology.

10 Toward a device that could stand

The preceding sections refute ACPCM as a *high-density* memory. They do not, by themselves, refute it as a *device* - provided it stops competing on the axis it loses and is rebuilt from the regimes where freezing water is *controllable* rather than chaotic. This section sketches that pivot in the same honest spirit as the rest of the paper: it remains speculative, but it is now *device-shaped* - a build recipe assembled from individually demonstrated components, a concrete operating point, and a fresh set of falsifiable tests. The key realization is that *directional* solidification of water is not a thought experiment but an industry: *freeze-casting* (ice-templating) routinely produces ordered, regular, repeatable ice architecture by controlling a freezing front [64]. “Freezing is too chaotic to record structure” is false in that regime; it is true only of the uncontrolled, deeply supercooled quench assumed in section 9.

10.1 Reframe the competition

ACPCM should not be pitched on density or energy, where it loses to DNA, 5D glass, and DRAM by the margins of section 6. It should be pitched on the axes where water is without peer: *substrate cost, abundance, non-toxicity, and recyclability*, deployed where cold is already free. Water is the cheapest write medium on Earth; it requires no synthesis (unlike DNA), is benign at any scale (unlike essentially every engineered medium), and is - uniquely - *infinitely re-writable at the substrate level*: melt it and the medium is pristine again. The target is the coldest, cheapest, most disposable tier of the storage hierarchy - the gravel of data storage, not the diamond.

10.2 Four moves that turn each obstacle into a design choice

Each load-bearing objection of section 9 maps to an established mechanism that neutralizes it, at the cost of abandoning the pure-water, high-density framing.

1. **Faithful commit** → **seed the scatterers**. Do not ask bare water to record the field. Seed the pod with acoustically responsive tracers and lock their field-assembled positions by freezing - the assemble-and-lock cycle already demonstrated with gelation [21, 22] (section 4.4). This retires Assumption 4 entirely; the recorded quantity is particle position, which *is* preserved through solidification.
2. **Dendritic overrun** → **stay below the instability**. The freeze front destroys fine structure only *above* the Mullins–Sekerka critical velocity [63]: below it, in a sufficiently steep thermal gradient, the solid–liquid interface is morphologically *stable* and advances as a planar (or regularly cellular) front. Freeze-casting exploits exactly this to template ordered lamellae [64]. Trade write speed for fidelity and the “front overruns the pattern” objection becomes a regime boundary to engineer around, not a law.
3. **No addressable grid** → **encode in the grain architecture**. Let the directional freeze itself carry the data: lamellar spacing, the grain-boundary map, and - the multi-level axis - the crystallographic *c*-axis orientation of each grain, steered by a *weak* (millitesla, not tesla) rotating magnetic bias during freezing, as in magnetic freeze-casting [65]. Orientation read by acoustic (or birefringence) anisotropy gives more than one bit per grain. The mechanism that section 9 called a confound - fields steering nucleation and growth - becomes the *signal*.
4. **Energy/throughput** → **co-opt ambient cold**. The $\sim 4 \mu\text{J/bit}$ freeze cost (section 6) is a cost only where heat must be actively pumped. Site the medium where cold is free and standing - polar archives, the cold aisle of a cryogenic data centre, or lunar permanently-shadowed regions at $\sim 40 \text{ K}$ - and the budget collapses toward the transducer and read electronics. Re-melt-to-rewrite, a drawback for RAM, is reframed as built-in secure erasure plus substrate recycling.

10.3 The honest operating point

This device is still acoustic-read-limited (section 6.2): $\sim 10^4\text{--}10^7 \text{ bits/mm}^3$, below every incumbent on volumetric density, and energy-competitive only where cold is ambient. We claim no density win and no energy win. The claim is narrower and, we think, defensible: a *non-empty niche* for an ultra-low-cost, non-toxic, substrate-recyclable, geologically stable cold-archive tier, built from components each demonstrated in isolation. Whether that niche is worth occupying is an economic question; that it is physically constructible to test is the point.

10.4 Device-mode predictions

The pivot earns its keep only if it, too, is falsifiable. It adds three tests to those of section 8.

Prediction 5 (Planar-front fidelity). Recording fidelity of a seeded pattern is high when the front advances *below* the Mullins–Sekerka critical velocity [63] and degrades sharply above it. *Falsified if* fidelity is independent of front velocity (which would mean the imprint is not a solidification-front effect at all).

Prediction 6 (Recyclability). The pod survives many melt–refreeze–rewrite cycles with no accumulating error floor. *Falsified if* dissolved-gas or solute build-up degrades readback within a few cycles - the failure mode that would make “infinitely re-writable” false in practice.

Prediction 7 (Orientation multi-level). A millitesla rotating bias during directional freezing sets per-grain *c*-axis orientation that survives the full freeze and is recoverable by acoustic anisotropy, yielding $> 1 \text{ bit/grain}$ [65]. *Falsified if* orientation does not survive solidification or cannot be read above chance.

Every block invoked here - acoustic assembly [21], planar-front directional solidification [63, 64], magnetic grain alignment [65], and acoustic-anisotropy read-out [56] - is established in isolation. The remaining conjecture is only that they *compose* into a working channel. That is the precise difference between a metaphor and a device, and it is a conjecture one can build a bench to test.

11 Conclusion

We have presented Acousto-Cymatic Phase-Change Memory as a structured thought experiment that transposes two proven storage principles - rapid phase-change writing (HAMR [3], PCM [9]) and microscopic isolation (CTF [5]) - onto a rapidly frozen aqueous medium, with an acoustic-holographic, neural-decoded read path that reuses otherwise-established techniques [12, 25]. Its value is not a feasibility claim but the discipline of the synthesis: by laying out an explicit epistemic ledger (table 2) and a falsifiable test program (section 8), the proposal makes clear exactly which single physical claim (Assumption 4) it lives or dies by, concentrates all genuinely speculative risk there, and quarantines refuted ideas as motivation only. And where the high-density framing fails, section 10 shows the idea need not collapse to metaphor: rebuilt from controlled directional solidification, seeded scatterers, and millitesla grain orientation - each established in isolation - it has a device-shaped, falsifiable operating point as a recyclable cold-archive medium, competing on cost and benignity rather than density. We offer the whole in that spirit: a provocation to be tested and, most likely, falsified - which is the useful kind.

Epistemic disclosure

This manuscript is a speculative “idea paper.” Where it invokes “water memory,” “structured/EZ water,” or W. Russell’s octave cosmology, it does so with explicit labels indicating these are contested, minority, or refuted positions, used as motivational or historical framing only; none carries the proposal’s logical weight. The established mechanisms cited (HAMR, charge-trap flash, phase-change memory, acoustic holography, deep-learning acoustic inversion, data sonification) are mainstream, and each is cited only to support its own subsystem; their combination into ACPCM is hypothetical and unvalidated. Nothing here should be read as a claim that the device works. All references were independently verified against Crossref, publisher, arXiv, or ADS records during preparation.

Acknowledgements

The synthesis draws on a reading list spanning storage engineering, acoustics, and the history of fringe science; the author thanks the open literature and explicitly credits the critical and skeptical sources [33, 34, 37] that establish the boundaries of accepted knowledge.

References

- [1] S. H. Charap, P.-L. Lu, and Y. He, “Thermal stability of recorded information at high densities,” *IEEE Trans. Magn.*, vol. 33, no. 1, pp. 978–983, Jan. 1997, doi: 10.1109/20.560142.
- [2] D. Weller and A. Moser, “Thermal effect limits in ultrahigh-density magnetic recording,” *IEEE Trans. Magn.*, vol. 35, no. 6, pp. 4423–4439, Nov. 1999, doi: 10.1109/20.809134.
- [3] M. H. Kryder, E. C. Gage, T. W. McDaniel, W. A. Challener, R. E. Rottmayer, G. Ju, Y.-T. Hsia, and M. F. Erden, “Heat assisted magnetic recording,” *Proc. IEEE*, vol. 96, no. 11, pp. 1810–1835, Nov. 2008, doi: 10.1109/JPROC.2008.2004315.
- [4] W. A. Challener *et al.*, “Heat-assisted magnetic recording by a near-field transducer with efficient optical energy transfer,” *Nat. Photonics*, vol. 3, no. 4, pp. 220–224, 2009, doi: 10.1038/nphoton.2009.26.
- [5] M. H. White, D. A. Adams, and J. Bu, “On the go with SONOS,” *IEEE Circuits Devices Mag.*, vol. 16, no. 4, pp. 22–31, Jul. 2000, doi: 10.1109/101.857747.
- [6] H. Tanaka *et al.*, “Bit cost scalable technology with punch and plug process for ultra high density flash memory,” in *2007 IEEE Symp. VLSI Technology, Tech. Dig.*, Jun. 2007, pp. 14–15, doi: 10.1109/VLSIT.2007.4339708.
- [7] A. Goda, “3-D NAND technology achievements and future scaling perspectives,” *IEEE Trans. Electron Devices*, vol. 67, no. 4, pp. 1373–1381, Apr. 2020, doi: 10.1109/TEDE.2020.2968079.

- [8] S. R. Ovshinsky, “Reversible electrical switching phenomena in disordered structures,” *Phys. Rev. Lett.*, vol. 21, no. 20, pp. 1450–1453, 1968, doi: 10.1103/PhysRevLett.21.1450.
- [9] M. Wuttig and N. Yamada, “Phase-change materials for rewriteable data storage,” *Nat. Mater.*, vol. 6, no. 11, pp. 824–832, 2007, doi: 10.1038/nmat2009.
- [10] S. Raoux, W. Wełnic, and D. Ielmini, “Phase change materials and their application to nonvolatile memories,” *Chem. Rev.*, vol. 110, no. 1, pp. 240–267, 2010, doi: 10.1021/cr900040x.
- [11] M. Faraday, “On a peculiar class of acoustical figures; and on certain forms assumed by groups of particles upon vibrating elastic surfaces,” *Phil. Trans. R. Soc. Lond.*, vol. 121, pp. 299–340, 1831, doi: 10.1098/rstl.1831.0018.
- [12] K. Melde, A. G. Mark, T. Qiu, and P. Fischer, “Holograms for acoustics,” *Nature*, vol. 537, no. 7621, pp. 518–522, 2016, doi: 10.1038/nature19755.
- [13] G. Memoli, M. Caleap, M. Asakawa, D. R. Sahoo, B. W. Drinkwater, and S. Subramanian, “Metamaterial bricks and quantization of meta-surfaces,” *Nat. Commun.*, vol. 8, art. 14608, 2017, doi: 10.1038/ncomms14608.
- [14] S. Jiménez-Gambín, N. Jiménez, J. M. Benlloch, and F. Camarena, “Holograms to focus arbitrary ultrasonic fields through the skull,” *Phys. Rev. Applied*, vol. 12, art. 014016, 2019, doi: 10.1103/PhysRevApplied.12.014016.
- [15] M. D. Brown, “Phase and amplitude modulation with acoustic holograms,” *Appl. Phys. Lett.*, vol. 115, no. 5, art. 053701, 2019, doi: 10.1063/1.5110673.
- [16] M. D. Brown, B. T. Cox, and B. E. Treeby, “Stackable acoustic holograms,” *Appl. Phys. Lett.*, vol. 116, no. 26, art. 261901, 2020, doi: 10.1063/5.0009829.
- [17] A. Marzo, S. A. Seah, B. W. Drinkwater, D. R. Sahoo, B. Long, and S. Subramanian, “Holographic acoustic elements for manipulation of levitated objects,” *Nat. Commun.*, vol. 6, art. 8661, 2015, doi: 10.1038/ncomms9661.
- [18] A. Marzo, M. Caleap, and B. W. Drinkwater, “Acoustic virtual vortices with tunable orbital angular momentum for trapping of Mie particles,” *Phys. Rev. Lett.*, vol. 120, art. 044301, 2018, doi: 10.1103/PhysRevLett.120.044301.
- [19] D. J. Collins, B. Morahan, J. Garcia-Bustos, C. Doerig, M. Plebanski, and A. Neild, “Two-dimensional single-cell patterning with one cell per well driven by surface acoustic waves,” *Nat. Commun.*, vol. 6, art. 8686, 2015, doi: 10.1038/ncomms9686.
- [20] J. Rufo, F. Cai, J. Friend, M. Wiklund, and T. J. Huang, “Acoustofluidics for biomedical applications,” *Nat. Rev. Methods Primers*, vol. 2, art. 30, 2022, doi: 10.1038/s43586-022-00109-7.
- [21] K. Melde, H. Kremer, M. Shi, S. Seneca, C. Frey, I. Niemann, S. Weißenbruch, S. Harrer, B. Maier, B. W. Drinkwater, and P. Fischer, “Compact holographic sound fields enable rapid one-step assembly of matter in 3D,” *Sci. Adv.*, vol. 9, no. 6, art. eadf6182, 2023, doi: 10.1126/sciadv.adf6182. [Acoustic assembly of particles/cells fixed permanently by gelation of the host medium.]
- [22] K. Olofsson, B. Hammarström, and M. Wiklund, “Acoustic focusing of beads and cells in hydrogel droplets,” *Sci. Rep.*, vol. 11, art. 8090, 2021, doi: 10.1038/s41598-021-86985-7. [Standing-wave focusing then cross-linking locks particle positions.]
- [23] International Electrotechnical Commission, *IEC 62127-1:2022, Ultrasonics - Hydrophones - Part 1: Measurement and characterization of medical ultrasonic fields*, 2nd ed. Geneva: IEC, 2022.
- [24] E. Martin and B. E. Treeby, “Investigation of the repeatability and reproducibility of hydrophone measurements of medical ultrasound fields,” *J. Acoust. Soc. Am.*, vol. 145, no. 3, pp. 1270–1282, 2019, doi: 10.1121/1.5093306.
- [25] Q. Lin, J. Wang, F. Cai, R. Zhang, D. Zhao, X. Xia, J. Wang, and H. Zheng, “A deep learning approach for the fast generation of acoustic holograms,” *J. Acoust. Soc. Am.*, vol. 149, no. 4, pp. 2312–2322, 2021, doi: 10.1121/10.0003959.
- [26] G. Kramer, B. Walker, T. Bonebright, P. Cook, J. H. Flowers, N. Miner, and J. Neuhoff, *Sonification Report: Status of the Field and Research Agenda*. Report prepared for the U.S. National Science Foundation by members of the Int. Community for Auditory Display (ICAD), 1999. [Online]. Available: <https://digitalcommons.unl.edu/psychfacpub/444/>

- [27] T. Hermann, A. Hunt, and J. G. Neuhoff, Eds., *The Sonification Handbook*. Berlin: Logos Verlag, 2011, ISBN 978-3-8325-2819-5. [Online]. Available: <https://sonification.de/handbook/>
- [28] T. Hermann, “Taxonomy and definitions for sonification and auditory display,” in *Proc. 14th Int. Conf. Auditory Display (ICAD)*, Paris, France, Jun. 2008.
- [29] J. W. Trayford, C. M. Harrison, R. C. Hinz, M. Kavanagh Blatt, S. Dougherty, and A. Girdhar, “Inspecting spectra with sound: proof-of-concept and extension to datacubes,” *RAS Tech. Instrum.*, vol. 2, no. 1, pp. 387–392, 2023, doi: 10.1093/rasti/rzad021.
- [30] W. Smith, “Interactive musical periodic table: sonification of visible element emission spectra,” in *Proc. 21st Sound and Music Computing Conf. (SMC)*, 2024.
- [31] M. L. Cowan, B. D. Bruner, N. Huse, J. R. Dwyer, B. Chugh, E. T. J. Nibbering, T. Elsaesser, and R. J. D. Miller, “Ultrafast memory loss and energy redistribution in the hydrogen bond network of liquid H₂O,” *Nature*, vol. 434, pp. 199–202, 2005, doi: 10.1038/nature03383.
- [32] S. Hameroff and R. Penrose, “Consciousness in the universe: a review of the ‘Orch OR’ theory,” *Phys. Life Rev.*, vol. 11, no. 1, pp. 39–78, 2014, doi: 10.1016/j.plrev.2013.08.002. [Contested hypothesis.]
- [33] M. Tegmark, “Importance of quantum decoherence in brain processes,” *Phys. Rev. E*, vol. 61, no. 4, pp. 4194–4206, 2000, doi: 10.1103/PhysRevE.61.4194. [Skeptical rebuttal to microtubule quantum coherence.]
- [34] D. C. Elton, P. D. Spencer, J. D. Riches, and E. D. Williams, “Exclusion zone phenomena in water - a critical review of experimental findings and theories,” *Int. J. Mol. Sci.*, vol. 21, no. 14, art. 5041, 2020, doi: 10.3390/ijms21145041. [Critical review of the “fourth phase” interpretation.]
- [35] G. H. Pollack, *The Fourth Phase of Water: Beyond Solid, Liquid, and Vapor*. Seattle, WA: Ebner & Sons, 2013, ISBN 978-0-9626895-4-3. [Self-published; contested minority hypothesis - cite with [34].]
- [36] E. Davenas, F. Beauvais, J. Amara, *et al.* (J. Benveniste), “Human basophil degranulation triggered by very dilute antiserum against IgE,” *Nature*, vol. 333, pp. 816–818, 1988, doi: 10.1038/333816a0. [Original “memory of water” claim; discredited - see [37].]
- [37] J. Maddox, J. Randi, and W. W. Stewart, “‘High-dilution’ experiments a delusion,” *Nature*, vol. 334, no. 6180, pp. 287–290, 1988, doi: 10.1038/334287a0. [On-site investigation: effect did not survive blinding.]
- [38] L. Montagnier, J. Aïssa, S. Ferris, J.-L. Montagnier, and C. Lavallée, “Electromagnetic signals are produced by aqueous nanostructures derived from bacterial DNA sequences,” *Interdiscip. Sci. Comput. Life Sci.*, vol. 1, no. 2, pp. 81–90, 2009, doi: 10.1007/s12539-009-0036-7. [Non-reproduced fringe revival.]
- [39] W. Russell, *The Universal One*. New York: privately published by the author, 1926. [Esoteric/occult cosmology; historical primary source for the octave/spiral metaphor only, not a chemistry reference. Facsimile: <https://archive.org/details/the-universal-one-1926-walter-russell>.]
- [40] J. Zhang, M. Gecevičius, M. Beresna, and P. G. Kazansky, “Seemingly unlimited lifetime data storage in nanostructured glass,” *Phys. Rev. Lett.*, vol. 112, no. 3, art. 033901, 2014, doi: 10.1103/PhysRevLett.112.033901.
- [41] Y. Lei, M. Sakakura, L. Wang, Y. Yu, H. Wang, G. Shayeganrad, and P. G. Kazansky, “High speed ultrafast laser anisotropic nanostructuring by energy deposition control via near-field enhancement,” *Optica*, vol. 8, no. 11, pp. 1365–1371, 2021, doi: 10.1364/OPTICA.433765.
- [42] G. M. Church, Y. Gao, and S. Kosuri, “Next-generation digital information storage in DNA,” *Science*, vol. 337, no. 6102, p. 1628, 2012, doi: 10.1126/science.1226355.
- [43] N. Goldman, P. Bertone, S. Chen, C. Dessimoz, E. M. LeProust, B. Sipos, and E. Birney, “Towards practical, high-capacity, low-maintenance information storage in synthesized DNA,” *Nature*, vol. 494, no. 7435, pp. 77–80, 2013, doi: 10.1038/nature11875.
- [44] Y. Erlich and D. Zielinski, “DNA Fountain enables a robust and efficient storage architecture,” *Science*, vol. 355, no. 6328, pp. 950–954, 2017, doi: 10.1126/science.aaj2038.
- [45] L. Organick *et al.*, “Random access in large-scale DNA data storage,” *Nat. Biotechnol.*, vol. 36, no. 3, pp. 242–248, 2018, doi: 10.1038/nbt.4079.

- [46] L. Ceze, J. Nivala, and K. Strauss, “Molecular digital data storage using DNA,” *Nat. Rev. Genet.*, vol. 20, no. 8, pp. 456–466, 2019, doi: 10.1038/s41576-019-0125-3.
- [47] J. F. Heanue, M. C. Bashaw, and L. Hesselink, “Volume holographic storage and retrieval of digital data,” *Science*, vol. 265, no. 5173, pp. 749–752, 1994, doi: 10.1126/science.265.5173.749.
- [48] H. J. Coufal, D. Psaltis, and G. T. Sincerbox, Eds., *Holographic Data Storage* (Springer Series in Optical Sciences). Berlin: Springer, 2000, ISBN 978-3-540-66691-2, doi: 10.1007/978-3-540-47864-5.
- [49] N. S. Osborne, “Heat of fusion of ice. A revision,” *J. Res. Natl. Bur. Stand.*, vol. 23, no. 6, pp. 643–646, 1939. [Classic calorimetric value $L_f = 333.5 \text{ J g}^{-1}$. Online: https://nvlpubs.nist.gov/nistpubs/jres/23/jresv23n6p643_A1b.pdf.]
- [50] W. Wagner and A. Pruß, “The IAPWS formulation 1995 for the thermodynamic properties of ordinary water substance for general and scientific use,” *J. Phys. Chem. Ref. Data*, vol. 31, no. 2, pp. 387–535, 2002, doi: 10.1063/1.1461829.
- [51] P. Brüggeller and E. Mayer, “Complete vitrification in pure liquid water and dilute aqueous solutions,” *Nature*, vol. 288, no. 5791, pp. 569–571, 1980, doi: 10.1038/288569a0.
- [52] J. Schmidt *et al.*, “Direct measurement of the critical cooling rate for the vitrification of water,” *Phys. Rev. Research*, vol. 7, art. 013095, 2025, doi: 10.1103/PhysRevResearch.7.013095.
- [53] P. Zhang, Z. Zhu, and D.-W. Sun, “Mechanism of ultrasound assisted nucleation during freezing and its application in food freezing process,” *Int. J. Food Prop.*, vol. 24, no. 1, pp. 68–88, 2021, doi: 10.1080/10942912.2020.1858862. [Cavitation-bubble nucleation and dendrite fragmentation; stochastic, not field-recording.]
- [54] A. E. Delgado, L. Zheng, and D.-W. Sun, “Influence of ultrasound on freezing rate of immersion-frozen apples,” and review of ice-nucleation control by ultrasonic, electric, and magnetic fields, *J. Food Eng.*, vol. 195, pp. 1–12, 2017, doi: 10.1016/j.jfoodeng.2016.09.011. [Review: external fields steer ice nucleation statistics.]
- [55] M. Saclier, R. Peczkalski, and J. Andrieu, “A theoretical model for ice primary nucleation induced by acoustic cavitation,” *Ultrason. Sonochem.*, vol. 17, no. 1, pp. 98–105, 2010, doi: 10.1016/j.ultsonch.2009.04.008. [Cavitation bubbles as the nucleation trigger.]
- [56] V. Anisimkin, V. Kolesov, A. Kuznetsova, E. Shamsutdinova, and I. Kuznetsova, “An analysis of the water-to-ice phase transition using acoustic plate waves,” *Sensors*, vol. 21, no. 3, art. 919, 2021, doi: 10.3390/s21030919. [MHz plate waves track the freezing transition acoustically.]
- [57] S. Biswas *et al.*, “Revealing the structure of snowflakes using acoustic levitation,” *Proc. Natl. Acad. Sci. USA*, 2025, doi: 10.1073/pnas.2502112122. [Levitated droplet freezes by inward dendritic growth; field holds, does not write the interior.]
- [58] A. L. Yarin, G. Brenn, O. Kastner, and C. Tropea, “Drying of acoustically levitated droplets of liquid–solid suspensions: evaporation and crust formation,” *Phys. Fluids*, vol. 14, no. 7, pp. 2289–2298, 2002, doi: 10.1063/1.1483308. [Suspended particles driven by the field; arrangement preserved into the dried/solid granule.]
- [59] M. Tagami, M. Hamai, I. Mogi, K. Watanabe, and M. Motokawa, “Solidification of levitating water in a gradient strong magnetic field,” *J. Cryst. Growth*, vol. 203, no. 4, pp. 594–598, 1999, doi: 10.1016/S0022-0248(99)00141-4. [Containerless freezing of diamagnetically levitated water.]
- [60] T. Naito, T. Suzuki, and Y. Ikezoe, “Diamagnetic levitation of water realized with a simple device consisting of ordinary permanent magnets,” *Appl. Phys. Lett.*, vol. 125, no. 26, art. 264102, 2024, doi: 10.1063/5.0241203. [~0.3 mm droplet levitated without a superconducting magnet.]
- [61] W.-K. Rhim, S. K. Chung, D. Barber, K. F. Man, G. Gutt, A. Rulison, and R. E. Spjut, “An electrostatic levitator for high-temperature containerless materials processing in 1-g,” *Rev. Sci. Instrum.*, vol. 64, no. 10, pp. 2961–2970, 1993, doi: 10.1063/1.1144475. [Charged-droplet containerless processing; deep undercooling.]

- [62] L. Otero, A. C. Rodríguez, M. Pérez-Mateos, and P. D. Sanz, “Effects of magnetic fields on freezing: application to biological products,” *Compr. Rev. Food Sci. Food Saf.*, vol. 15, no. 3, pp. 646–667, 2016, doi: 10.1111/1541-4337.12202. [Review: magnetic-field effects on ice nucleation and crystal size.]
- [63] W. W. Mullins and R. F. Sekerka, “Stability of a planar interface during solidification of a dilute binary alloy,” *J. Appl. Phys.*, vol. 35, no. 11, pp. 444–451, 1964, doi: 10.1063/1.1713333. [Planar front is stable below a critical growth velocity; above it, cellular/dendritic.]
- [64] S. Deville, E. Saiz, R. K. Nalla, and A. P. Tomsia, “Freezing as a path to build complex composites,” *Science*, vol. 311, no. 5760, pp. 515–518, 2006, doi: 10.1126/science.1120937. [Freeze-casting: directional solidification of water yields ordered, regular ice architecture at scale.]
- [65] M. M. Porter, M. Yeh, J. Strawson, T. Goehring, S. Lujan, P. Siripasopsotorn, M. A. Meyers, and J. McKittrick, “Magnetic freeze casting inspired by nature,” *Mater. Sci. Eng. A*, vol. 556, pp. 741–750, 2012, doi: 10.1016/j.msea.2012.07.058. [A weak rotating magnetic field orients grains during directional freezing - long-range order at millitesla scale.]
- [66] A. A. Shibkov, Yu. I. Golovin, M. A. Zheltov, A. A. Korolev, and A. A. Leonov, “Morphology diagram of nonequilibrium patterns of ice crystals growing in supercooled water,” *Physica A*, vol. 319, pp. 65–79, 2003, doi: 10.1016/S0378-4371(02)01517-0.
- [67] Y. Xu, N. G. Petrik, R. S. Smith, B. D. Kay, and G. A. Kimmel, “Growth rate of crystalline ice and the diffusivity of supercooled water from 126 to 262 K,” *Proc. Natl. Acad. Sci. USA*, vol. 113, no. 52, pp. 14921–14925, 2016, doi: 10.1073/pnas.1611395114.
- [68] M. Horowitz, “Computing’s energy problem (and what we can do about it),” in *2014 IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2014, pp. 10–14, doi: 10.1109/ISSCC.2014.6757323.