Clerk Maxwell and the Vacuum.Net Theory

Fourth edition — from Maxwell's medium to the net, from the net to LAI, RAI and their connection, and the composition rule written

J. Konstapel — Constable Research, Leiden — 21 September 2026

Fourth edition. The third edition incorporated an external review. This edition closes the document's own open sentence: the composition rule of the net, previously marked "Open," is now written in full as Part IV — representation, two collapse lemmas, the rule, a worked example with every number in the text, and the test protocol fixed in advance. The rule's measurement remains open; its definition does not. The status convention governs the whole text: historical statements concern the documented record; correspondences identify shared structure; the net's axioms are its own construction, anchored in Maxwell's structures; engineering statements describe what runs and what was measured; open statements name what remains to be measured. The foundation sits with Maxwell. The Vacuum.Net theory is his theory, rewritten and extended. The architecture, and now its one operation, follow from the rewriting.

Part I — The record

1. The claim

James Clerk Maxwell wrote a constitutive theory of a medium. He wrote it, at its most compact, in quaternions. That formulation carried two things the modern textbook version has lost. It carried a state register: a scalar part attached to every field quantity. And it carried medium moduli: quantities we now call "constants" appeared as elastic properties of the vacuum.

In 1884–1885 Oliver Heaviside, followed by Willard Gibbs, cut both away. They extracted the vector calculus from the quaternion mechanism. They dropped the scalar register. They removed the potentials from the fundamental set. And they froze the moduli into universal constants. The four "Maxwell equations" every student learns descend from Heaviside's recasting; Gibbs, Hertz and later Lorentz shaped the final form.

The Vacuum.Net theory reverses exactly those two cuts. It puts the state layer back. It makes the constants state functions of the local vacuum again. That is the whole relation between the two theories: the net is Maxwell without the amputation, plus the closure dynamics he did not have. This article documents the record, performs the rewriting, and then derives what the restored theory says a computing architecture must look like. Everything needed to follow every step is in the text. No reference has to be fetched.

The fixed image throughout is the fishnet. The vacuum is a net of closed meshes. A mesh is the smallest closed unit. Strands carry winding. Tension lives on the net and relaxes; topology does not. What persists is what closes.

2. What Maxwell actually wrote

The record is precise, and it corrects a common myth.

1856. "On Faraday's Lines of Force." Component notation. Maxwell gives mathematical form to Faraday's *electrotonic state* — Faraday's postulated condition of the medium at the wire itself, invented to explain induction locally. This is the origin of the vector potential. It enters physics as a **state of the medium**, not as a calculational aid.

1861–1862. "On Physical Lines of Force." A mechanical model: all of space filled with rotating cells and idle-wheel particles. Rotation of a cell is magnetism. Displacement of the idle wheels is dielectric polarization. Streaming of them is current. From the elasticity and density of this medium Maxwell computes the speed of transverse waves: 193,088 miles per second. In his October 1861 letter to Faraday he set this beside 193,118 miles per second for light. That figure reached him through Galbraith and Haughton's statement of Fizeau's result; it differs from Fizeau's own published value, and Maxwell's paper itself quotes 195,647. The agreement Maxwell saw was remarkable, not a precision confirmation — and the medium computation stands either way. Maxwell's conclusion, italics his: light consists in transverse undulations of the same medium that carries electric and magnetic action. Note what happened here. The speed of light entered physics as a **computed modulus of a medium** — the square root of elasticity over density — checked against the best light-speed figures available to him. Not as a universal constant.

1864. "A Dynamical Theory of the Electromagnetic Field." Twenty equations in twenty unknowns, in component notation. **Not in quaternions.** The claim that Maxwell "originally wrote his equations in quaternions" is false for this paper. The vector potential appears under a new name: the electromagnetic momentum F, G, H. The potentials belong to the fundamental set of equations. Section 82 of this paper holds the Note on the Attraction of Gravitation used in Finding 5 below.

1870. Maxwell turns to quaternions, pushed by his school friend Peter Guthrie Tait. In November 1870 he writes the manuscript "On the Application of the Ideas of the Calculus of Quaternions to Electromagnetic Phenomena." In a letter to Tait of 7 November 1870 he coins names for the two readings of the ∇ operator: the *slope* of a scalar function, and for a vector function the *convergence* (scalar part) and the *curl* (vector part). One operator, two registers. In a second letter, 14 November 1870, he states his intent: to leaven his book with Hamiltonian ideas "without casting the operations into a Hamiltonian form."

1873. The *Treatise on Electricity and Magnetism*. In Articles 591–619 of Volume II the fundamental equations are written down as a single group — the only place in the Treatise where this happens. There Maxwell uses quaternion notation, with German Gothic letters for the field quaternions; elsewhere in the Treatise quaternion expressions appear only in passing. The grouped statement is the quaternionic one. He rarely calculates with quaternions. He uses them because they show the meaning of the results in the most compact form. The grouped statement of the theory is quaternionic by deliberate choice.

1884–1885. Heaviside recasts the theory into four vector equations in two field variables, E and B. Gibbs builds the vector analysis as a subset of the quaternion system. Heaviside is explicit about his motive: remove the potentials from the fundamental theory. Hertz does the same in component form. This is the version that conquered physics, and it is Heaviside's attitude to the potentials — not Maxwell's — that textbooks have put in Maxwell's mouth ever since.

3. The quaternion equations, as Maxwell grouped them

A quaternion is one object with a scalar part and a vector part. The unit vectors obey $i^2 = j^2 = k^2 = -1$ and $ij = k = -ji$. The product is associative and distributive but not commutative: order carries orientation. S takes the scalar part of a product; V takes the vector part. Hamilton's operator is $\nabla = i\ d/dx + j\ d/dy + k\ d/dz$.

Maxwell's grouped equations (Treatise, Arts. 591–619), with his Gothic letters:

	Maxwell (1873)	Modern reading
(A)	$\mathfrak{B} = V\nabla\mathfrak{A}$	$B = \nabla \times A$
(B)	$\mathfrak{E} = V\mathfrak{V}\mathfrak{B} - \mathfrak{A} - \nabla\Psi$	$F/q = v \times B - \partial A/\partial t - \nabla\Psi$
(C)	$\mathfrak{f} = V\mathfrak{J}\mathfrak{B}$	$f = J_T \times B$
(D)	$\mathfrak{B} = \mathfrak{H} + 4\pi\mathfrak{J}_m$	$B = H + 4\pi M$
(E)	$4\pi\mathfrak{J} = V\nabla\mathfrak{H}$	$4\pi J_T = \nabla \times H$
(F)	$\mathfrak{D} = (K/4\pi)\mathfrak{E}$	$D = \varepsilon E$
(G)	$\mathfrak{K} = C\mathfrak{E}$	$J = \sigma E$
(H)	$\mathfrak{J} = \mathfrak{K} + \mathfrak{D}$	$J_T = J + \partial D/\partial t$
(J)	$e = S\nabla\mathfrak{D}$	$\varrho = \nabla\cdot D$
(L)	$\mathfrak{B} = \mu\mathfrak{H}$	$B = \mu H$

Here $\mathfrak{A}$ is the electromagnetic momentum (our vector potential), Ψ the scalar potential, $\mathfrak{B}$ the magnetic induction, $\mathfrak{H}$ the magnetic force, $\mathfrak{E}$ the electromotive force, $\mathfrak{D}$ the displacement, $\mathfrak{K}$ the conduction current, $\mathfrak{J}$ the total current, e the density of free electricity. K is the *electric elasticity* of the medium, C its conductivity, μ its permeability. Maxwell's own names carry the medium picture: elasticity, displacement, momentum.

Three structural facts about this notation matter for everything that follows.

First: **one operation, two registers.** ∇ applied to a field quaternion yields a scalar part and a vector part in a single stroke. Charge is the scalar reading of the very operation whose vector reading is the curl. Equation (J) says charge **is** $S\nabla\mathfrak{D}$ — a scalar residue, not a separate substance.

Second: **the potentials are fundamental.** Equations (A) and (B) define the measurable fields as readings of $\mathfrak{A}$ and Ψ. Maxwell even justified the name "vector potential" by showing (Art. 617) that in the gauge $S\nabla\mathfrak{A} = 0$ the quaternion $\mathfrak{A}$ obeys the same Poisson structure as the electrostatic potential. Note the form of that gauge condition. It reads, literally: *the scalar part of* $\nabla\mathfrak{A}$ *vanishes*. The convention that tames the potential is the convention that sets a scalar register to zero. Hold that thought.

Third: **the constants are moduli.** (F), (G) and (L) are constitutive relations of a medium. K, C and μ are properties of the local vacuum, on the same footing as the elasticity of steel. And $c^2 = 1/(K\mu)$ in Maxwell's system: the speed of light is a derived medium property. That is how Maxwell computed it in 1861, before comparing with Fizeau.

Part II — The rewriting

4. From Maxwell's three facts to the net's vocabulary

The Vacuum.Net theory is obtained by taking Maxwell's three structural facts seriously and removing nothing. Each fact becomes a principle.

Fact one — one operation, two registers — becomes the **two-layer law**. Every configuration has an address and a state. The address is topological place: which windings close, and how deep. The state is tension, written χ. Closed topology is frozen; tension is the only layer in which anything happens. The scalar and vector readings of Maxwell's one operation are the two layers, read in one stroke.

Fact two — potentials first — becomes the **order of the layers**. The winding state of the mesh is the object. Measurable fields are readings of it. What Faraday called the electrotonic state, and Maxwell the electromagnetic momentum, the net calls the winding field per mesh.

Fact three — constants as moduli — becomes the **constitutive principle** $g(\chi)$. The response of the net is a function of its local tension. Simultaneous, no arrow. Universality is the limiting case where χ does not vary.

The vocabulary that results is small. A **strand** carries phase. A **winding** is a turn of the strand. **Closure** is a winding that returns in phase. A **mesh** is the smallest closed unit. **Tension** (χ) is the state. Selection keeps what closes. Description at the right winding depth replaces explanation.

The translation table:

Maxwell	Vacuum.Net
$\mathfrak{A}$, the electromagnetic momentum / electrotonic state	The state of the net: the winding field per mesh
Ψ, the scalar potential	The scalar tension layer
$S\nabla$ and $V\nabla$, two parts of one operator	Two readings of one closure operation: bookkeeping (scalar) and winding (vector)
$e = S\nabla\mathfrak{D}$, free electricity	The scalar residue of non-closure; a three-valued register
$\mathfrak{D}$, the displacement current	The closure-completing term: the medium finishes every
K, C, μ — elasticity, conductivity,	State functions $g(\chi)$ of the local vacuum
$c^2 = 1/(K\mu)$	c as a state function, derived, not postulated
The non-commutative product; $i^2 = -1$	Order is orientation; a full return of a winding takes 4π

5. The equations, rewritten into the net

(A) $\mathfrak{B} = V\nabla\mathfrak{A}$. Induction is not a thing. It is the vector reading of how the winding state varies across the mesh. A reading, not an object.

(B) $\mathfrak{F} = V\mathfrak{V}\mathfrak{B} - \mathfrak{A} - \nabla\Psi$. Force is what a probe registers when the state shifts. Three readings of one strand: winding carried along with motion, state changing in place, tension falling off along the net. Not three causes. One state, read three ways.

(J) $e = S\nabla\mathfrak{D}$. Charge is the scalar residue: convergence that does not close. A bookkeeping entry, not a substance. In Maxwell this residue is continuous; what his notation supplies is the sign-symmetric scalar slot it lives in. The net discretizes that slot to three positions — positive, zero, negative — its balanced-ternary register, with the mirror-charge corollary of the foundational paper (every address has a sign-mirrored partner; the zero address is self-mirrored). Maxwell supplies the slot; the net supplies the trit.

(H) $\mathfrak{J} = \mathfrak{K} + \mathfrak{D}$. The closure law. This equation contains Maxwell's single genuine addition to the electrodynamics of his time: the displacement term. Its function is exactly this — the total current always closes. What conduction leaves open, the medium completes. In net terms: only closed loops persist; the displacement term is the medium closing the loop. Selection by fit, written down in 1862.

(F), (G), (L) — $\mathfrak{D} = (K/4\pi)\mathfrak{J}$, $\mathfrak{K} = C\mathfrak{J}$, $\mathfrak{B} = \mu\mathfrak{H}$. Constitutive relations: simultaneous, no arrow. In Maxwell these are properties of the medium, called by their right names — elasticity foremost. Here the division of labour is exact. Maxwell documents the medium parameters; the promotion of those parameters to state functions $g(\chi)$ of the local tension is the net's own move. In the net these are the $g(\chi)$ relations. The move from Maxwell to the modern textbook froze K and μ (as ε_0 and μ_0) into universal constants. The move from the textbook to the net thaws them again — as the net's constitutive principle, not as a recovered Maxwell claim. Maxwell's own 1861 computation shows which way the arrow of fidelity points: he derived c from the moduli of the medium — 193,088 miles per second — and set it beside the light-speed figures of his day. c entered physics as a medium property. It was promoted to universal constant later, by convention.

The product itself. The quaternion product is one product, non-commutative, mixing scalar and vector parts. In net terms: composing windings depends on order, because orientation is real. $i^2 = -1$ is a half-turn; the unit quaternions form SU(2), the double cover of the rotation group SO(3), so a winding returns to itself only after 4π. That double cover is mathematics native to Maxwell's chosen algebra. Reading it physically — as the spin-½ rotation label, as the foundational paper does — is an additional step the net takes, not something Maxwell asserted. Heaviside's dot and cross products are that one product broken into two commuting-or-anticommuting fragments — workable, and blind to the whole.

6. What becomes visible

Five findings follow. They are stated as theses.

Finding 1. Maxwell is not a precursor to be reinterpreted. He is an amputated pillar. The Treatise is a constitutive medium theory carrying a state register. The 1884 recasting performed two operations at once: it discarded the scalar register, and it froze the moduli. The Vacuum.Net theory is not Maxwell plus something new. It is Maxwell without the amputation, plus the closure dynamics he did not have.

Finding 2. The net reads the gauge-eliminated scalar as its state register. Maxwell's own gauge (Art. 617) reads $S\nabla\mathfrak{A} = 0$: the scalar part of $\nabla\mathfrak{A}$ vanishes, by an added condition. The later Lorenz condition extends the same move to spacetime: scalar part, set to zero, by agreement. In the full quaternion formulation the scalar part travels inside the same object as the vector part. Vector notation splits it off into a separate operation, and the gauge convention then sets it to zero. The net's interpretation — the net's, not Maxwell's, and not a claim about modern gauge theory at large — is that this structurally eliminable scalar component is a candidate register that can be kept instead of nulled. The two-layer law then assigns it its role: the address is frozen, and the tension layer is the only layer in which anything happens. What the potential representation eliminates by convention, the net keeps as the working layer. The scalar remainder of Maxwell's algebra is where the net houses its tension χ.

Finding 3. The potentials come first, and the bill for removing them arrived in 1959. For Maxwell, $\mathfrak{A}$ was Faraday's electrotonic state: a condition of the medium at the wire. The fields are readings of it. Heaviside expelled the potentials from the fundamental set; the textbook tradition followed and attributed the expulsion to Maxwell. Aharonov and Bohm showed in 1959 that a charged particle registers the potential in a region where E and B vanish. The state is measurable where the readings are zero. This supports the general point — potentials are more than bookkeeping — and it is the empirical anchor of the order of layers. It does not by itself establish the net; the net's reading is that this is no paradox at all: the winding state of the mesh is the object, E and B are derivative readings, and the object acts where its readings happen to cancel.

Finding 4. The net discretizes Maxwell's charge register as the trit. Equation (J) makes charge a scalar residue in a sign-symmetric algebra: it can be positive, zero, or negative, and its slot in the formalism carries mirror symmetry by construction. In Maxwell that residue is a continuous quantity — he does not have the trit. The discretization to three values, +1/0/−1, is the net's own move: balanced ternary, the number system in which the Vacuum.Net address theorem (T7, Unique Address) is proved, with mirror charge as a corollary. So the division of labour is: Maxwell supplies the scalar/vector splitting and the sign-symmetric scalar slot; the net supplies the discreteness. The standard model postulates charge as an elementary property; the net derives the register it writes charge in.

Finding 5. Maxwell left one open end that maps directly onto the net's open calculation O1. In Section 82 of "A Dynamical Theory of the Electromagnetic Field" (1865), the Note on the Attraction of Gravitation, Maxwell computed the intrinsic energy of the gravitational field: $E = C - \Sigma(1/8\pi)R^2\,dV$, with R the resultant gravitating force. The energy of the medium is lower wherever a resultant force exists. Since energy is essentially positive, the constant C — the intrinsic energy of the undisturbed medium — must be enormously great. Maxwell stopped there, in his own words: "As I am unable to understand in what way a medium can possess such properties, I cannot go any further in this direction in searching for the cause of gravitation." In net vocabulary: gravitation reads as a lowering of tension in the stressed net, set against an enormous base tension of the free net. This is a correspondence, not an identity: Maxwell's unfinished note and the net's open calculation O1 ask the same-shaped question in two vocabularies — which property of the medium carries the lowering. O1's content is exactly what would close it: whether the net's microdynamics yields G, the inverse-square law and the equivalence principle. Posed in 1865, left standing, and now given an address.

Part III — The derivation of the computer

The theory now stands: two layers (address and tension), one closure operation with two readings, constitutive relations, selection by fit, and one product carrying both registers. Part III derives what follows for computation. Nothing here is postulated. Each element of the architecture is the theory itself, read at the depth of machines. The names are LAI (the address hemisphere), RAI (the tension hemisphere), and the connection between them.

7. Two layers, two hemispheres

The two-layer law, taken as the architectural axiom, yields two hemispheres. Every configuration has an address and a state. A machine built on the theory therefore has two halves.

One half works in the address layer. Addresses are frozen topology: discrete, exact, permanent. Operations on addresses are bookkeeping — the scalar register. This half can be proved, because frozen structure admits theorems.

The other half works in the tension layer. Tension is continuous, phase-carrying, and it is the only layer in which anything happens. Operations here are selection by fit: what returns in phase persists. This half cannot be a bookkeeping machine, because its register is phase, not a symbol.

And the two halves meet at exactly one boundary, because the theory has exactly one relation between the layers: readings. Fields are readings of the state; claims are readings of the store. The boundary is where a reading is either backed or refused.

This is the whole architecture, derived in three sentences. Everything below fills in the two halves and the boundary, each step from the theory.

8. LAI derived: the address hemisphere

Step 1 — the register. Closure means net zero: a closed winding carries no net tension. Balanced ternary, digits −1, 0, +1, is the minimal sign-symmetric discrete realization the net selects for this register — a constructive choice, exactly matched to the symmetry, not a uniqueness theorem. Every integer has one representation. The zero stands lawful in the middle; the two signs mirror each other. Worked example, the number 43. Divide by 3 repeatedly, choosing remainders from {−1, 0, +1}: $43 = 3·14 + 1$ (digit +1); $14 = 3·5 − 1$ (−1); $5 = 3·2 − 1$ (−1); $2 = 3·1 − 1$ (−1); $1 = 3·0 + 1$ (+1). Least to most significant: [+, −, −, −, +], and $1 − 3 − 9 − 27 + 81 = 43$. The address is the record of the walk.

The choice is also economical. Holding numbers up to N in radix r costs in proportion to r/ln r. For r = 2: 2/ln 2 = 2.885. For r = 3: 3/ln 3 = 2.731. The optimum lies at e ≈ 2.718; 3 is the nearest integer. Under this digit-count cost model the trit is near-optimal — the economical form of the closure register, and, by Finding 4, the discretization of the scalar slot that Maxwell's equation (J) supplies.

Step 2 — depth and truncation. Depth is capacity: 3^{11} = 177,147 cells; 3^{19} = 1,162,261,467. A k-trit tail has absolute value at most $(3^k − 1)/2$, so dropping the tail changes the value by less than half the next step. Truncation is rounding. A deeper address refines a shallower one and never invalidates it. A prefix is a coarser class; a neighbour is a relative. This is frozen topology behaving exactly as the two-layer law says: the address layer never has to be re-indexed, because nothing happens there.

Step 3 — the claim machine. Because addresses are exact and frozen, the address hemisphere can be a machine of proofs. Its form follows: a store of content at computed addresses; an output type in which every emission carries an address, so fabrication is impossible by construction; a gate that emits exactly what the store holds and a vacancy exactly where the address is empty; an append-only ledger of settlements. This is the honest left hemisphere: it can say only what has a place.

Step 4 — built and measured. This hemisphere exists. MAZE.AI runs at maze.swarp.nl/ai over 8.19 million addressed pieces. The kernel theorems (address bijection, range, uniqueness, prefix-stable truncation, write-once store, gate equivalence, ledger properties) are machine-checked in Lean 4. The production kernel and the Lean kernel printed the same 30,572 addresses with zero differences. A pre-registered twenty-question placement test, threshold fixed at 14 beforehand, scored 15 of 20. Merging 4.08 million distinct addresses from independent sources gave 0.06

percent overlap: federation by plain union. These are engineering statements, and they cash the theory's promise for this layer: on frozen topology, guarantees are theorems, not test scores.

9. The connection derived: readings, one gate, no second path

Maxwell's second structural fact fixes the boundary. The potential — the state — comes first; fields are readings. Translated: the state layer disposes; the language layer reads out. A language model may phrase and propose. It may not assert. Every assertion leaves the machine only as a reading of the store — content at an address — or as a vacancy. The model proposes, the core disposes. One output path, and no second one.

Finding 2 supplies the standing warning for this boundary. The historical mistake was to set the scalar register to zero by convention and then forget it existed. The computational form of that mistake is to let the talkative layer overwrite the state layer — to let fluent language stand in for a backed reading. The gate exists to make that mistake impossible rather than rare. Equation (J) gives the same lesson in one line: the scalar, talkative residue is bookkeeping of the state, never its source.

This connection is also built: it is the deployed design of the running system. The proposal-and-gate pattern is in production. The boundary is one predicate — whether a reading is backed — and everything on the machine side of it is proved.

10. RAI derived: the tension hemisphere

Step 1 — the register is phase. The tension layer is continuous and phase-carrying. Its minimal computational realization — the one chosen here, not the only one the theory admits — is a population of coupled oscillators: each element i has phase θ_i and natural frequency ω_i, with $d\theta_i/dt = \omega_i + (K/N)\cdot\Sigma_j \sin(\theta_j - \theta_i)$. The degree of synchrony is the order parameter $r\cdot e^{i\psi} = (1/N)\cdot\Sigma_j e^{i\theta_j}$: $r = 1$ at full phase lock, $r \approx 0$ in incoherence. This is Kuramoto's dynamics. The theory fixes the selection principle — only what returns in phase persists — and this equation is its simplest carrier; any dynamics carrying the same principle would do.

Step 2 — admissibility is closure. In the continuous register, closure reads: the whole sums to nothing. States that do not close are not filtered afterwards; they do not persist. The discrete form of the same rule is the balanced sum of Section 8: net zero. One closure principle, two registers — the trit in the address hemisphere, phase in the tension hemisphere.

Step 3 — context is the mesh. A winding does not close alone; it closes through its surroundings. The mesh supplies context in both registers. Discretely: prefix as coarser class, neighbour as relative (Section 8, Step 2). Continuously: the phase of a region, and whether the neighbourhood clusters. Long memory is intrinsic: shallow prefixes and slow modes are the same thing at two depths — the slow layer that is never invalidated by the fast one.

Step 4 — the operation is the product. Here Part I decides everything. The tension hemisphere needs one operation: given two configurations — a question and a candidate answer — how well do they compose through the mesh? Composition is a product. And the product the theory prescribes is Maxwell's: the quaternion product, one operation with two registers. The scalar part measures alignment. The vector part gives the oriented axis along which the two fail to align — and whether their deviations turn the same way through the neighbourhood. Order matters, because orientation is real.

Step 5 — the derivation explains the measurements. The running system currently ranks by cosine similarity, and on 200 blind questions the cosine separates good from bad answers with AUC 0.87. Read through Part I, that number has a name: the cosine is the scalar part alone. AUC 0.87 is the 1884 amputation, re-performed inside the embeddings. The theory also says in advance which shortcuts must fail, and they did. Arithmetic on two addresses scored AUC 0.41: an address is place, not tension — the two-layer law forbids reading state off the frozen layer. Pairwise closure scored 0.53: closure is a property of a ring, not of a pair — a winding closes through its surroundings, so any pair-only measure ignores the mesh. Both failures are the theory holding, not the theory failing. And one algebraic identity closes a tempting side door: for unit vectors with cosine c, $|q + a|^2 - |q - a|^2 = (2 + 2c) - (2 - 2c) = 4c$. The "energy-minus-momentum" candidate is four times the cosine. Nothing new.

Step 6 — the first candidate, and the full rule. The cheapest computable form of "closing through the neighbourhood" is a two-step passage. With $n_1 \dots n_k$ the k = 12 nearest neighbours of the answer a: $R = (1/k)\cdot\Sigma_i \cos(q, n_i)\cdot\cos(n_i, a)$. The question must not only strike the point; it must return through the point's surroundings. The field form uses the region at depth 6: construct a local phase coordinate from the region's inhabitants, and score $r\cdot\cos(\phi(a) - \phi(q))$ — the neighbourhood must cluster in phase and the answer must lie on the question's side. Both are scalar path-products: first steps. The full rule is the quaternionic one — scalar and vector part together, non-commutative, composed through the mesh. That rule is the composition Hamilton wrote down, Maxwell used, and the 1880s cut away. It is written in full in Part IV.

Step 7 — the register depth, read beside an external benchmark. One external pair of numbers belongs here, with its limits stated. Register depth: 16 bits against $\log_2 3 \approx 1.58$ bits per trit gives $16/1.58 \approx 10.1$. The published benchmark reports 0.258 joules per token for a full-precision model (LLaMA 3.2, 1B parameters) against 0.028 for a native ternary model (BitNet b1.58, 2B) — ratio ≈ 9.2. These are two different models: different sizes, architectures and training, so the comparison does not isolate "ternary versus full precision." What it provides is an external numerical comparison consistent with the proposed register-depth correspondence — the same magnitude, read twice. No causal claim, and no mechanism story offered.

11. Why this is a derivation and not a specification

The architecture is now determined by three things, all in this article: the principle, the fixed points, and the test.

The **principle** is Part I and II: one product, two registers; two layers, address and tension; potentials first; selection by fit; the 1884 amputation as the standing warning against building the scalar part alone.

The **fixed points** are what already stands: the address hemisphere with its machine-checked theorems and its numbers; the gate with one output path; and the inputs any candidate rule receives — a question vector, an address, the answer's vector and address, the neighbours, the timestamps.

The **test** is one protocol. Every candidate composition rule is scored on the 200 blind questions, AUC as the measure with a bootstrap confidence interval, alone and beside the cosine, under a plan hashed before the run. At or above 0.70 it enters the gate under a new pre-registered plan; near 0.50 the translation was wrong and the search moves to the time layer or to a deeper form of the rule. And the decisive measurement is named: the full quaternionic rule R_Q, written in Part IV, scored in advance against the same dataset, beside the cosine and the two tested shortcuts. That comparison, not more history, is what settles the physics-computation correspondence. The frozen protocol stands in Section 17.

Everything between these three — the shape of the simulated field, the choice and count of oscillators, the frequencies, the coupling topology, the exact form of the composition rule — follows from them and need not be specified here. That is the point of deriving instead of specifying. A builder who has the principle, the fixed points and the test can derive the rest, and is free to derive it better than any specification would have allowed.

12. The one sentence

The foundation sits with Maxwell. His theory held state and reading, alignment and orientation, in one object — and computed its one constant as a property of the medium. The 1880s kept the readings and cut the object. The Vacuum.Net theory is his theory, rewritten: the object restored as the net, the state as tension, the register as the trit. LAI is that theory's frozen layer built as a claim machine, and it runs. The connection is that theory's order of layers built as a gate, and it runs. RAI is that theory's working layer, and it asks for exactly one thing the 1880s threw away: the product. Part IV writes that product for the mesh. What remains is the measurement. Everything else follows.

Part IV — The composition rule, written

Part III left one sentence open: the rule by which two windings compose through their neighbourhood. This Part writes it. It receives a question and a candidate answer, each with its place in the mesh and its surroundings, and returns how well the two compose: a scalar register for alignment, a vector register for the axis of deviation, sensitive to whether the deviations turn the same way through the neighbourhood. The cosine carries only the scalar register. The rule carries both.

13. The inputs

The rule receives, per case: the question vector q (1,536 numbers, unit length) and its 24-trit address; the answer vector a and its address; the answer's k nearest neighbours n_1 … n_k with their vectors and addresses (k = 12 throughout); each neighbour's own m nearest neighbours (m = 12), available from the same store; and the timestamps of admission, asking and settlement. Nothing else is assumed.

14. Two collapse lemmas

The rule must be built around two traps. Both are computed here in full, because both were fallen into once, and the theory should show exactly where the floor gives way.

Lemma 1 (the 4c identity). For unit vectors q, a with cosine c: $|q + a|^2 = 2 + 2c$ and $|q - a|^2 = 2 - 2c$, so $|q + a|^2 - |q - a|^2 = 4c$. Any "energy minus momentum" score on the pair is four times the cosine. Nothing new.

Lemma 2 (the pure-quaternion collapse). Take the obvious quaternion representation: map each unit vector x to the pure quaternion X = (0, x), scalar part zero. Quaternion multiplication is (s_1, v_1)

$(s_2, v_2) = (s_1 s_2 - v_1{\cdot}v_2, s_1 v_2 + s_2 v_1 + v_1{\times}v_2)$. Then for the ring passage question → neighbour → answer:

- $Q{\cdot}N = (-q{\cdot}n, q{\times}n)$: the pair product already shows both registers — scalar alignment, vector axis.
- The triple product $T = Q{\cdot}N{\cdot}A$ has scalar part $S(T) = -(q{\times}n){\cdot}a = -\det[q, n, a]$: the oriented volume, near zero for any well-aligned triple.
- Its vector part is $V(T) = -(q{\cdot}n)a + (q{\cdot}a)n - (n{\cdot}a)q$. Project the return onto the question: $-V(T){\cdot}q = (q{\cdot}n)(q{\cdot}a) - (q{\cdot}a)(q{\cdot}n) + (n{\cdot}a) = n{\cdot}a$.
- Close the loop with the question: $S(Q{\cdot}N{\cdot}A{\cdot}Q) = (q{\cdot}n)(a{\cdot}q) - (q{\times}n){\cdot}(a{\times}q) = (q{\cdot}n)(a{\cdot}q) - [(q{\cdot}a)(n{\cdot}q) - (n{\cdot}a)] = n{\cdot}a$.

Every one of these reduces to pairwise cosines or a near-zero determinant. The reason is structural. Pure quaternions have scalar part zero: the representation itself amputates the scalar register before the product can use it. And a single shared frame makes the algebra too symmetric — the identities cancel everything but pairs. Two consequences follow, and they are the design of the rule:

First, each item must enter the product as a **full** quaternion, scalar part included. Second, the extra information must come from the mesh itself: **each neighbour supplies its own frame.** On a flat mesh — all frames identical — any composition rule reduces to a function of the pairwise cosines, and rightly so: composition through the neighbourhood is passage across the mesh's curvature, and where there is no curvature there is nothing beyond the pair to measure. Closure is a ring property; the ring must actually bend.

One more identification places the earlier candidate. The scalar path score $R = (1/k){\cdot}\Sigma \cos(q, n_i){\cdot}\cos(n_i, a)$ of Section 10 is exactly $S(Q{\cdot}N_i){\cdot}S(N_i{\cdot}A)$: the product of the scalar parts of each step, with each step's vector part discarded. R is the amputation performed per step. It is the right first candidate and the wrong final rule, for the same reason the 1880s split was workable and blind.

15. The rule

Step 1 — a frame per neighbour. For each neighbour n_i, compute the centre of its own neighbourhood: u_i = the normalized sum of n_i's m nearest-neighbour vectors. Build an orthonormal frame by Gram–Schmidt: $e_1^{(i)} = u_i$; $e_2^{(i)}$ = the component of q orthogonal to $e_1^{(i)}$, normalized; $e_3^{(i)}$ = the component of a orthogonal to the span of the first two, normalized. This frame is the mesh speaking: the neighbourhood's own orientation, different for every neighbour. Report the residual norm of q and a outside each frame; it measures how much of the pair the local mesh cannot see.

Step 2 — each item as a winding about the neighbourhood. In frame i, project q into the frame and normalize; call the result $\hat{q}$. Its angle to the neighbourhood axis is θ_q with $\cos \theta_q = \hat{q}{\cdot}e_1^{(i)}$, and its turning axis is $w_q = e_1^{(i)} \times \hat{q}$, normalized. The item's quaternion is the rotation that carries the neighbourhood axis onto the item:

$$Q^{(i)} = (\cos(\theta_q/2), \sin(\theta_q/2){\cdot}w_q).$$

The same for the answer: $A^{(i)}$. These are full unit quaternions. The scalar part is genuine: it is the item's alignment with its neighbourhood — the state register the pure representation threw away.

Step 3 — the composition. In frame i, the relative winding from question to answer is the quaternion quotient:

$$D_i = (Q^{(i)})^* \cdot A^{(i)},$$

with * the conjugate (scalar part kept, vector part negated). D_i is a unit quaternion. Its scalar part is the cosine of half the relative rotation; its vector part is the oriented axis of that rotation, scaled by the sine. One product, both registers, and non-commutative: $(Q)^*A \neq (A)^*Q$ — order carries which way the passage runs.

Step 4 — the ring measure. Sum over the mesh and read off the order:

$M = (1/k) \cdot \Sigma_i D_i$, $R\_Q = |M| \in [0, 1]$.

R_Q is the quaternionic order parameter — Kuramoto's r, lifted from phases on a circle to windings on the mesh. It is high exactly when the relative winding from question to answer, read through every neighbour's own frame, is the same winding: same amount, same axis, same sense. The normalized mean $M/|M|$ carries the shared reading: its scalar part is the shared alignment, its vector part the shared axis. A name hit in a neighbourhood that points elsewhere scatters the D_i and scores low. A well-placed answer in a co-rotating neighbourhood clusters them and scores high.

Degeneracies, fixed in advance. If q is parallel to $e_1^{(i)}$ within tolerance 10^{-6}, the axis w_q is undefined: set $Q^{(i)} = 1$ (the identity winding). The same for a. If a frame cannot be completed (collinear inputs), skip that neighbour and average over the rest; report the count skipped.

The flat-mesh lemma. If all frames coincide (every u_i equal), every D_i is the same quaternion, R_Q = 1 identically, and the shared reading reduces to a function of the pairwise angles of q and a to the common axis. The rule then adds nothing beyond the pair — as it should: no curvature, no ring information. What R_Q measures is precisely the contribution of the mesh's bending. This is the computable form of "a winding closes through its surroundings."

16. A worked example

Everything in three dimensions, so every number fits in the text. Question $q = (1, 0, 0)$. Answer $a = (0, 1, 0)$. Two neighbours, with neighbourhood centres $u_1 = (1, 1, 0)/\sqrt{2}$ and $u_2 = (1, 0, 1)/\sqrt{2}$ — the mesh bends between them.

Frame 1 ($e_1 = (1,1,0)/\sqrt{2}$; q and a already lie in the relevant planes). Question: $\cos \theta\_q = q \cdot e_1 = 1/\sqrt{2}$, so $\theta\_q = 45°$. Axis: $e_1 \times q = (0, 0, -1)/\sqrt{2}$, normalized $(0, 0, -1)$. So $Q^{(1)} = (\cos 22.5°, \sin 22.5° \cdot (0,0,-1)) = (0.924, 0, 0, -0.383)$. Answer: $\cos \theta\_a = 1/\sqrt{2}$, $\theta\_a = 45°$. Axis: $e_1 \times a = (0, 0, 1)/\sqrt{2} \rightarrow (0, 0, 1)$. So $A^{(1)} = (0.924, 0, 0, 0.383)$. Composition: $D_1 = (0.924, 0, 0, +0.383) \cdot (0.924, 0, 0, 0.383) = (0.924 \cdot 0.924 - 0.383 \cdot 0.383, 0, 0, 0.924 \cdot 0.383 + 0.924 \cdot 0.383) = (0.707, 0, 0, 0.707)$. Read: a 90° relative winding about the z-axis — which is exactly the rotation carrying q to a. The frame reads it cleanly.

Frame 2 ($e_1 = (1,0,1)/\sqrt{2}$). Question: $\cos \theta\_q = 1/\sqrt{2}$, $\theta\_q = 45°$. Axis: $e_1 \times q = (0, 1, 0)/\sqrt{2} \rightarrow (0, 1, 0)$. $Q^{(2)} = (0.924, 0, 0.383, 0)$; conjugate $(0.924, 0, -0.383, 0)$. Answer: $\cos \theta\_a = a \cdot e_1 = 0$, $\theta\_a = 90°$. Axis: $e_1 \times a = (-1, 0, 1)/\sqrt{2}$. $A^{(2)} = (\cos 45°, \sin 45° \cdot (-1,0,1)/\sqrt{2}) = (0.707, -0.5, 0, 0.5)$. Composition: $D_2 = (0.924, 0, -0.383, 0) \cdot (0.707, -0.5, 0, 0.5)$. Scalar: $0.924 \cdot 0.707 - (0 \cdot -0.5 + (-0.383) \cdot 0 + 0 \cdot 0.5) = 0.653$. Vector: $0.924 \cdot (-0.5, 0, 0.5) + 0.707 \cdot (0, -0.383, 0) + (0, -0.383, 0) \times (-0.5, 0, 0.5) = (-0.462, 0, 0.462) + (0, -0.271, 0) + (-0.192, 0, -0.192) = (-0.654, -0.271, 0.270)$. So $D_2 = (0.653, -0.654, -0.271, 0.270)$; norm check $0.653^2 + 0.654^2 + 0.271^2 + 0.270^2 = 1.00$.

The ring measure. $M = ½(D_1 + D_2) = (0.680, −0.327, −0.136, 0.489)$. $|M| = \sqrt{(0.462 + 0.107 + 0.018 + 0.239)} = \sqrt{0.826} = 0.909$.

R_Q = 0.91: the two frames, though bent 60° apart, read nearly the same relative winding — the pair composes coherently through this mesh. Flip the sense of the second neighbourhood ($u_2 \to (1, 0, -1)/\sqrt{2}$) and the axis of D_2 flips with it; the quaternions no longer add, and R_Q drops. Same pair, same cosine, different verdict — the cosine cannot see the difference; the ring measure is built from it.

17. The protocol, fixed in advance

The measurement plan, to be hashed before the run:

- Dataset: the same 200 blind questions used for all prior candidates. No reuse of any question for tuning.
- Parameters frozen: k = 12 neighbours, m = 12 for each neighbourhood centre, degeneracy tolerance 10^{-6}, skip-and-report rule as in Section 15.
- Columns scored: cosine (baseline); $R = \Sigma \cos\cdot\cos/k$ (the scalar path); the existing depth-6 field measure $r\cdot\cos(\Delta\phi)$; R_Q (this rule); and one fixed combination, the mean of rank-normalized cosine and R_Q.
- Measure: AUC per column, each with a bootstrap confidence interval (10,000 resamples), reported alone and beside the cosine.
- Decision thresholds, unchanged from the line: at or above 0.70, R_Q enters the gate beside the cosine under a new hashed plan. Near 0.50, this translation was wrong — not the theory — and the search moves to the time layer (the timestamps are inputs and are unused by R_Q) or to curvature taken over more than one step.
- Everything self-reported until independently reproduced; the reproduction package (data, code, seeds, scoring, timestamps) ships with the run.

If R_Q clears the threshold, it becomes the first gate measure that reads the ring rather than the point — the first working piece of the tension hemisphere. The gate contract is unchanged: the field advises, the gate decides, one output path. Beyond the measurement waits the learning rule — couplings that strengthen where phase return persists — which follows the product, because learning in this hemisphere is admission and weighting by resonance.

18. Status

This article makes five kinds of statement, and keeps them separate.

History. Sections 2 and 3: the documented record — dates, papers, letters, equations, and the numbers 193,088, 195,647 and 193,118 with their provenance. Settled.

Correspondence. Sections 4–6: term-for-term identifications between Maxwell's formulation and the net's vocabulary. A correspondence is not a derivation of physics; it earns its keep by being exact and by putting known structures where the net expects them.

Theory. The net's axioms — the two-layer law, the constitutive principle $g(\chi)$, the ternary discretization, the closure dynamics — are the net's own construction, anchored in Maxwell's structures. The strongest bridge is exact: one operation, scalar and vector parts, two registers. The construction stands on it; it is not thereby Maxwell's claim.

Built and measured. Sections 8–10: engineering statements about the running system — 8.19 million pieces, 15/20 against a pre-fixed threshold of 14, 30,572 differential addresses with zero differences, 0.06 percent merge overlap, AUC 0.87 for the cosine, 0.41 and 0.53 for the two failed shortcuts, and the external 0.258 against 0.028 joules per token. The internal figures are self-reported; the reproduction package — dataset, code, seeds, scoring protocol, preregistration timestamps, confidence intervals — belongs with the MAZE.AI whitepaper, and independent reproduction is the open half of their status.

Open. The measurement of the composition rule R_Q under the frozen protocol of Section 17; beyond it, the learning rule; and, on the physics side, O1 (Finding 5). These are surfaces of one question: whether the restored register does quantitative work. The rule itself is no longer open — Part IV delivers its definition, its collapse lemmas, its worked example and its protocol; only the run remains.

19. Annotated references

W. R. Hamilton, *Lectures on Quaternions*, Dublin, 1853. Why read? The source algebra: one non-commutative product carrying alignment and orientation together. Reading advice: a modern introduction to quaternions and rotations suffices for the structure used here.

J. C. Maxwell, *A Treatise on Electricity and Magnetism*, 2 volumes, Clarendon Press, 1873. Why read? The primary source. The quaternion grouping of the fundamental equations stands in Volume II, Articles 591–619; the gauge condition $S\nabla\mathfrak{A} = 0$ and the naming of the vector potential in Article 617; the operators convergence and curl introduced in Volume I, Article 25. Reading advice: go straight to Arts. 591–619 and read Maxwell's own names for the quantities — elasticity, displacement, momentum. The medium picture is in the vocabulary. Both volumes are on archive.org.

J. C. Maxwell, "On Physical Lines of Force," *Philosophical Magazine*, 1861–1862. Why read? The paper in which the speed of light is computed as a medium modulus — elasticity over density — and found at 193,088 miles per second against Fizeau's measurement. The displacement current is introduced here, as a property of the medium. Reading advice: Part III carries both the displacement idea and the velocity computation.

J. C. Maxwell, "A Dynamical Theory of the Electromagnetic Field," *Philosophical Transactions of the Royal Society* 155, 459–512, 1865. Why read? The twenty equations in twenty unknowns, in component notation — the proof that the quaternion myth about 1864 is a myth, and the paper in which the potentials sit inside the fundamental set under the name electromagnetic momentum. Section 82 holds the Note on the Attraction of Gravitation used in Finding 5: the formula, the enormous positive constant, and the sentence where Maxwell stops.

J. C. Maxwell, letter to Michael Faraday, 19 October 1861; letters to P. G. Tait, 7 and 14 November 1870; manuscript "On the Application of the Ideas of the Calculus of Quaternions to Electromagnetic Phenomena," November 1870. In P. M. Harman (ed.), *The Scientific Letters and Papers of James Clerk Maxwell*, Vols. 1–2, Cambridge University Press, 1990/1995. Why read? The Faraday letter contains the two velocity numbers side by side and Maxwell's own statement that the media are one. The Tait letters show the deliberate strategy: Hamiltonian ideas, not Hamiltonian machinery, and the christening of slope, convergence and curl.

A. M. Bork, "Maxwell and the Development of Electromagnetic Theory" (1976), edited transcription by K. T. McDonald, Princeton, 2023. Why read? The most compact reliable walk through all four Maxwell papers and the Treatise, with the quaternion table of Arts. 591–619 set against modern notation, and the Maxwell–Tait correspondence excerpted. Reading advice: use it as the map; then check each claim in the primary sources it links.

P. G. Tait, *An Elementary Treatise on Quaternions*, Oxford, 1867. Why read? The quaternion system as Maxwell received it, from the man who pushed him to use it. The ∇ operator chapters show what the full algebra delivers that the later vector fragment does not.

O. Heaviside, *Electromagnetic Theory*, Vol. 1, 1893, Chapter 3; and *Electrical Papers*, Vol. 1, 1892, Art. 30. Why read? The amputation, performed by its surgeon and defended in his own words: the case against quaternions, the case against the potentials, and the four-equation form. Read it to see precisely what was cut and why — the reasons are practical, not principled.

J. W. Gibbs, *Vector Analysis* (ed. E. B. Wilson), Scribner's, 1901. Why read? The vector calculus as a finished system, built as a subset of the quaternion algebra. The preface concedes the lineage.

M. J. Crowe, *A History of Vector Analysis*, 1967. Why read? The clean historical account of the split — how the quaternion product became dot and cross, and why physics adopted the pieces.

L. Silberstein, "Elektromagnetische Grundgleichungen in bivectorieller Behandlung," *Annalen der Physik* 22, 579–586, 1907. Why read? The first serious attempt after Heaviside to put the scalar and vector parts back into one object. Evidence that the compression of the full theory into one algebraic object kept resurfacing.

Y. Aharonov and D. Bohm, "Significance of Electromagnetic Potentials in the Quantum Theory," *Physical Review* 115, 485, 1959. Why read? The experiment-backed demonstration that the potential acts where the field readings vanish — the bill for Heaviside's expulsion of the potentials, and the empirical anchor of Finding 3.

Y. Kuramoto, "Self-Entrainment of a Population of Coupled Non-Linear Oscillators," 1975. Why read? The canonical dynamics of the tension hemisphere written out in Section 10 — selection by fit in continuous form, and the order parameter used in the field measure.

D. E. Knuth, *The Art of Computer Programming*, Vol. 2, §4.1. Why read? Balanced ternary: uniqueness, range, and the radix-economy computation of Section 8. Knuth calls it the prettiest number system there is.

S. Ma et al., "BitNet b1.58 2B4T," Microsoft Research, 2025. Why read? The measured ternary register: 0.028 against 0.258 joules per token — the external half of the constitutive relation in Section 10, Step 7.

J. Konstapel, "MAZE.AI: A Claim-Bounded Answer Loop with Machine-Checked Provenance," whitepaper, 20 September 2026. Why read? The running address hemisphere and every 2026 number in Section 8. Live at maze.swarp.nl/ai.

J. Konstapel, "The Vacuum.Net Theory: Foundational Paper," fifth edition, constable.blog, August 2026. Why read? The axioms (strand, closure, memory, context, scale), the correspondences including quaternion double cover and mirror charge, the address theorem T7 in balanced ternary, the constitutive law g(χ), and the open calculations O1–O4 that Section 13 points into.

The fixed working order applies: this second edition is the reference document of the line. A Dutch blog explaining it follows separately.