V1.0.0

The Speed of Light as the Tracing Capacity of a Distributed-Estimation Vacuum: Causal Bandwidth, Stability Thresholds, and Divergence Times

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

The speed of light $c = 2.998 \times 10^{8}\ \mathrm{m/s}$ is measured, codified in the SI, and used daily, yet no established framework explains why the causal speed is finite or what physical function finiteness performs beyond kinematics. This paper develops a control-theoretic answer: the vacuum is modeled as a distributed estimation network whose nodes track unstable field dynamics through nearest-neighbor communication on a lattice of spacing a, so that c is the network's tracing capacity — the fastest instability the network can track without its self-estimation error diverging. From the lattice dispersion $\omega(k) = (2v/a)\,|\sin(ka/2)|$ we derive the exact causal bound $|v_g| \le v$; from the per-update error recursion we derive the stability threshold $\lambda_{\mathrm{crit}} = -\ln(1-g)\,c/a$ on the trackable instability rate at correction gain g. Under the Planck-spacing assumption $a = \ell_P = 1.616 \times 10^{-35}\ \mathrm{m}$ and $g = 1/2$, the threshold is $\lambda_{\mathrm{crit}} = 1.286 \times 10^{43}\ \mathrm{s}^{-1}$. An instability running at 1.5 times this threshold overruns a 10^6-fold error tolerance within $2.15 \times 10^{-42}\ \mathrm{s}$, about 40 update intervals. The model reproduces the causal bandwidth $B = 5.905 \times 10^{42}\ \mathrm{Hz}$, the congestion delay $3.336 \times 10^{-12}\ \mathrm{s/m}$ at $\rho/\rho_c = 10^{-3}$, and a device-scale cutoff $4.772 \times 10^{10}\ \mathrm{Hz}$ for millimetre geometry, and satisfies its own Nyquist consistency check with margin $1/\pi$. The result converts a kinematic axiom into a falsifiable stability condition and specifies the analogue lattice simulation that would confirm or refute it.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). The Speed of Light as the Tracing Capacity of a Distributed-Estimation Vacuum: Causal Bandwidth, Stability Thresholds, and Divergence Times. QNFO Preprint v1.0.0.

1. Introduction

Physicists working on quantum gravity, cosmologists confronting the horizon problem, and metrologists who codify c in the SI all face the same explanatory gap: the speed of light is measured, fixed, and used daily, but no existing framework explains why the causal speed is finite at all, or what physical role finiteness plays beyond kinematics [10]. The gap has practical consequences. Variable-speed-of-light cosmologies exist as responses to the horizon problem — the observation that the cosmic microwave background is isotropic to a few parts in 10^5 over regions that were causally disconnected at recombination [3] — but in those treatments c is modulated by an external mechanism, such as probe-brane motion in a bulk spacetime, rather than derived from the substrate's own operational constraints [3]. Relativity itself can be extended to accommodate a variable c with minimal changes, which shows that the kinematic structure is flexible; what remains missing is a reason for the specific finite value [2].

The gap blocks a specific design decision. If c is the tracing capacity of a distributed control network — the maximum instability rate the network can track before its self-estimation error diverges — then Lorentz invariance stops being a postulate and becomes a stability condition, and the value of c becomes a calculable quantity tied to the network's spacing and update protocol. That derivation would change what counts as evidence in quantum-gravity research: discrete-substrate programs that already derive a light-cone structure from graph spectral properties [8] would gain a control-

theoretic criterion for why the bound must hold, and engineering programs that build distributed feedback systems with explicit stability budgets — such as accelerator networks in which a beam position must be held stable to the order of $50\ \mathrm{nm}$ for an intra-train collision feedback loop to stay in capture range [4] — would share a formal vocabulary with foundational physics.

This paper formalizes that conjecture for a concrete model. The vacuum is modeled as a lattice of estimator nodes, each tracking unstable field modes through nearest-neighbor communication at a finite rate. The central claim is that the causal speed equals the network's tracing capacity: exceeding it makes the estimation error diverge, and the resulting stability boundary reproduces a Lorentz-invariant causal cone. The contributions are:

- Derivation of the exact causal bound $|v_g| \le v$ for the lattice dispersion $\omega(k) = (2v/a)\,|\sin(ka/2)|$, given in Section 4.1.
- Computation of the Planck-spaced causal bandwidth $\omega_{\max} = 3.710 \times 10^{43}\ \mathrm{s}^{-1}$ and $B = 5.905 \times 10^{42}\ \mathrm{Hz}$, with a Nyquist consistency check showing the bandwidth uses a fraction $1/\pi$ of the available sampling headroom, derived in Section 4.2.
- Derivation of the per-update error recursion and the exact stability threshold $\lambda_{\mathrm{crit}}(g) = -\ln(1-g)\,c/a$ on the trackable instability rate, with a sensitivity evaluation over gains $g \in \{0.1, 0.5, 0.9\}$, given in Sections 4.3 and 4.4.
- A divergence-timescale result: an instability at $1.5\lambda_{\mathrm{crit}}$ overruns a 10^6-fold error tolerance within 2.15×10^{-42} s, approximately 40 update intervals, derived in Section 4.5.
- Independent recomputation of the congestion delay $3.336 \times 10^{-12}\ \mathrm{s/m}$ at $\rho/\rho_c = 10^{-3}$ and the device-scale cutoff 4.772×10^{10} Hz for millimetre geometry, confirming internal consistency against the results of [10], derived in Sections 4.6 and 4.7.
- A formal statement of the causal cone $|\Delta x| \le c\,\Delta t$ with its isotropy assumptions, given in Section 4.8, and falsification conditions with a proposed minimal lattice simulation, stated in Sections 6 and 7.

The derivation could also bear on problems outside fundamental physics wherever a network must track dynamics faster than its communication allows: the same threshold structure governs any distributed estimator of an unstable process, and the capacity–fidelity trade it exposes parallels the payload-versus-robustness trade measured in learned steganographic channels [5].

2. Background and Related Work

The closest prior derivation of the speed of light from a discrete substrate is the Bruhat–Tits-tree program. That work maps the timeless adjacency operator of the Bruhat–Tits tree to a time-evolution operator via the Page–Wootters mechanism, derives a causal light-cone structure from Lieb–Robinson bounds, and concludes that the dimensionless upper bound on the speed of light is the Lieb–Robinson velocity derived from the graph's spectral radius, with physical dimensions supplied by the observer's thermodynamic modular Hamiltonian; c is thereby a derived, macroscopic, epistemic property of combinatorial geometry rather than a fundamental axiom [8]. The present paper shares the conclusion that c is derived from a network's spectral and structural properties, but supplies a complementary mechanism: where the tree program uses Lieb–Robinson propagation bounds on a p-adic geometry, this paper uses the stability of a distributed estimator tracking unstable dynamics, which produces an explicit inequality on the trackable instability rate rather than a propagation bound alone.

The regularization framing of the same conjecture is developed in the distributed-estimation-vacuum paper. That work derives an exact causal bound $|v_g| \le v$ for the lattice dispersion $\omega(k) = (2v/a)\,|\sin(ka/2)|$, evaluates the causal bandwidth $\omega_{\max} = 2c/\ell_P = 3.710 \times 10^{43}\ \mathrm{s}^{-1}$ and $B = \omega_{\max}/2\pi = 5.905 \times 10^{42}$ Hz for Planck spacing, shows that an assumed first-order response suppresses above-cutoff modes at -20 dB per decade — the quantitative signature of early stopping — states a stability window $\omega_s = \omega_c\sqrt{\Gamma_0/\gamma_0}$, and derives congestion delays $\Delta\tau = 3.336 \times 10^{-12}$ s

per metre at $\rho/\rho_c = 10^{-3}$ with a device-scale cutoff for millimetre-scale geometry [10]. Crucially, its discussion states what the model does not establish: that the physical vacuum is such a lattice, that the response function has the assumed form, or that the congestion variable maps onto the metric tensor; the conjecture is formalized and made falsifiable, not confirmed [10]. The present paper extends that program along the axis it leaves open: it adds the estimation-error dynamics themselves, derives the divergence threshold on the tracked instability, and quantifies how fast coherence is lost when the threshold is exceeded.

Variable-speed-of-light cosmology provides the physical motivation and the contrast case. In the brane-induced realization, the observed four-dimensional universe is a probe brane moving in a higher-dimensional bulk toward a black hole produced by a stack of branes, and the speed of light on the probe brane decreases in time as the brane approaches the horizon; a rotating black-hole background prevents the brane from hitting the horizon, bulk expansion provides effective friction, and freezeout followed by virialization stabilizes the orbit so that the speed of light tends to a constant at late times [3]. The introduction of that work frames the horizon problem quantitatively: the cosmic microwave background is isotropic to a few parts in 10^5 over regions causally disconnected at recombination, which standard cosmology cannot explain through causal microphysical processes [3]. That treatment derives a time-varying c from the substrate's dynamics but leaves the finiteness of c itself unexplained; the present model addresses finiteness directly as a stability requirement.

On the kinematic side, the extended-relativity work shows that special and general relativity survive a variable light speed with minimal changes. It extends the Lorentz transform and the Minkowski line element under the single assumption that all moving observers at the same spacetime point measure the same light speed, even though that speed may vary in time, and introduces the differential construct $d\hat{T} = c\,dT$ together with a rest mass m that varies so as to keep the rest energy $\hat{m} = mc^2$ constant, generating covariant velocity, energy-momentum and force vectors, with the field equation of general relativity extended to the variable-c case [2]. This matters for the present paper because it establishes the testable kinematic consequence: if the vacuum network's tracing capacity drifted, the extended transform of [2] describes how relativity would respond, so the model's prediction of a constant c is a substantive claim about network stationarity rather than a kinematic necessity.

The discrete-substrate ontology underlying the lattice assumption is articulated in the syntactic-generation framework. That work avoids positing any material substrate and takes as its single foundational element a syntactic token — the drawing of a boundary separating inside from outside — with a syntactic void acting as the identity element against which boundaries hold structural meaning, and with strictly recursive combination rules that generate infinite complexity from a minimal alphabet; its reduction law of idempotence makes two adjacent identical boundaries equivalent to one [9]. This supplies the conceptual warrant for modeling the vacuum as discrete nodes with local relations rather than a continuum: the present paper's assumption A1, that field dynamics live on a lattice of spacing a, is the physical reading of a boundary-combination structure. The epistemic-interface framing of the "All Things Are One" synthesis — in which the experienced interface is generated by an observer subgraph interacting with a static relational network, and model success does not guarantee correspondence to ontology — reinforces the paper's own honesty about what a formal model can and cannot establish [11].

Three experimental and computational works supply the operational vocabulary the model borrows. The structured-light work demonstrates that a topological index alone can govern propagation: an input beam carrying no initial spin and no modal asymmetry develops spin-dependent redistribution during propagation exclusively from the Pancharatnam topological index ℓ_p, which determines the distinct paraxial modal families into which the two spin components evolve, deterministically rather than as a generic consequence of paraxial propagation; its experimental section validates the prediction through Stokes polarimetry of a HeNe beam shaped by a liquid-crystal spatial light modulator across three propagation planes [1]. The relevance is structural: discrete indices selecting which modal families propagate is exactly the behavior the lattice model exhibits when the tracing capacity selects which field modes remain

causally coherent. The tomography work shows that a full optical state — spatial, spectral, temporal and polarisation — can be reconstructed from intensity-only measurements through a single-mode fibre with a single-pixel detector, including for mutually incoherent fields that spectrally overlap, by leveraging internal references so that no auxiliary light source is needed; its principle section frames measurement as projection of an unknown density matrix through a sequence of analyser states given by eigenvectors of Pauli matrices encoded on a spatial light modulator [7]. This is the measurement-side analogue of the estimator model: a band-limited observer reconstructing a field state from projected, finite-rate samples. The light-shift spectroscopy work provides the same lesson in atomic physics: the spatially varying ac-Stark shift of a dipole trap maps the radial location of each atom into a unique frequency shift, so a single-shot snapshot of the light-shift-broadened spectrum infers atom number, temperature, trap depth and transition linewidth with high independence between parameters, validated in a hollow-core fibre loaded with a magneto-optically trapped rubidium ensemble [6]. In all three cases a continuous physical distribution is recovered through a discrete, capacity-limited channel — the operating assumption of the vacuum estimator taken here.

Finally, the accelerator-engineering literature shows that distributed control networks with explicit stability budgets are built and operated at the relevant level of rigor. The ILC Technology Network's final-doublet work package requires the vertical beam fluctuation at the interaction-point quadrupole to be stable to the order of $50\ \mathrm{nm}$ to stay within the capture range of the intra-train collision feedback, evaluated through a cryostat vibration experiment with a superfluid-helium service line routed through radiation shielding; its conclusions report that responsible institutes are assigned for all work packages, with completion anticipated by April 2027 under current funding, and that the work packages carry relevance to other accelerator projects including the FCC and synchrotron light sources [4]. The steganography work completes the capacity vocabulary: StegNet hides a same-size image with a decoding rate of 98.2%, corresponding to 23.57 bits per pixel, by changing only 0.76% of the cover coefficients, explicitly trading accuracy for capacity and noting that the embedded image need not even be visually similar to the cover; its survey of traditional methods — image-domain, transform-domain and file-format-based, the last being fragile because any alteration crashes the hidden payload — quantifies how channel capacity and robustness trade against each other, the same trade the vacuum estimator faces between tracking bandwidth and estimation fidelity [5].

3. Methods

The paper's method is formal modeling with explicit derivation and computation; no experiment, simulation on hardware, or data collection is performed. The model, assumptions, derivation procedure and computation are as follows.

Model. The vacuum is modeled as a network of estimator nodes on a one-dimensional nearest-neighbor lattice with spacing a. Each node holds an estimate $\hat{\phi}_i(t)$ of an underlying unstable field ϕ and updates it by exchanging information with its nearest neighbors at a finite communication speed v. Two descriptions of the same network are used and must be consistent: a wave description, in which the lattice supports modes with dispersion relation

$$\omega(k) = \frac{2v}{a}\left|\sin\!\left(\frac{ka}{2}\right)\right|, \tag{1}$$

and an estimation description, in which each node's error $e_i = \hat{\phi}_i - \phi$ evolves under a per-update recursion derived in Section 4.3. The wave description yields the causal bound and the bandwidth; the estimation description yields the stability threshold and the divergence time. Their consistency is checked in Section 4.2.

Assumptions.

- **A1 (Lattice structure).** Field dynamics live on a nearest-neighbor lattice of spacing a, taken as the Planck length $\ell_P = 1.616 \times 10^{-35}\ \mathrm{m}$ for numerical evaluation. This is a projection: physical realizations would employ an effective spacing determined by the analogue system.

- **A2 (Communication speed).** All nodes exchange information at the same speed $v = c = 2.998 \times 10^8$ m/s, the defined SI value.
- **A3 (First-order response).** Above the lattice cutoff the network's transfer function behaves as a single-pole low-pass filter, $|H(\omega)| = 1/\sqrt{1 + (\omega/\omega_c)^2}$, giving -20 dB attenuation per decade.
- **A4 (Update protocol).** Each node applies a fractional correction of gain g per update interval $\Delta t = a/v$, and the tracked instability grows at rate λ between updates.
- **A5 (Congestion model).** The additional propagation delay per metre depends linearly on a dimensionless density ratio ρ/ρ_c, with proportionality constant $\alpha = 1/c$, so that $\Delta\tau = \alpha\,\rho/\rho_c$. We take $\alpha = 1/c$ under the convention that at unit density ratio the propagation delay equals the light-travel time per metre; this reproduces the value used in [10].

Procedure. All results are obtained analytically: the causal bound from differentiating Eq. (1); the bandwidth from evaluating Eq. (1) at the zone edge; the stability threshold from the error recursion of A4; the divergence time from iterating that recursion; and the congestion delay and device cutoff by direct substitution. Every arithmetic step is shown in Section 4. No numerical simulation is executed in this work; the derived quantities are intended to guide future analogue-simulation experiments.

4. Analysis

4.1 Causal bound

Differentiating Eq. (1) with respect to k gives the group velocity

$$v_g(k) = \frac{\partial \omega}{\partial k} = v \cos\!\Big(\frac{ka}{2}\Big). \tag{2}$$

Since $|\cos\theta| \le 1$ for any real θ,

$$|v_g(k)| = \Big|v \cos\!\Big(\frac{ka}{2}\Big)\Big| \le v, \tag{3}$$

with the bound saturated at $k = 0$, where $\cos(0) = 1$ and $v_g = v$. The lattice therefore enforces the exact causal bound $|v_g| \le v$ for every wavevector in the Brillouin zone.

4.2 Causal bandwidth and Nyquist consistency

The maximum of Eq. (1) over the Brillouin zone is attained at the zone edge, where $|\sin(ka/2)| = 1$:

$$\omega_{\max} = \frac{2v}{a}. \tag{4}$$

Inserting the inputs $v = c = 2.998 \times 10^8$ m/s (A2) and $a = \ell_P = 1.616 \times 10^{-35}$ m (A1):

$$\omega_{\max} = \frac{2 \times 2.998 \times 10^8 \ \mathrm{m/s}}{1.616 \times 10^{-35} \ \mathrm{m}} = \frac{5.996 \times 10^8}{1.616 \times 10^{-35}} \ \mathrm{s^{-1}} = 3.710 \times 10^{43} \ \mathrm{s^{-1}}. \tag{5}$$

The corresponding ordinary frequency bandwidth is

$$B = \frac{\omega_{\max}}{2\pi} = \frac{3.710 \times 10^{43}}{6.2832} \ \mathrm{Hz} = 5.905 \times 10^{42} \ \mathrm{Hz}. \tag{6}$$

Nyquist consistency check. The update interval is $\Delta t = a/v$, so the sampling angular frequency of the network is $\omega_s = 2\pi/\Delta t = 2\pi v/a = \pi\,\omega_{\max}$. The band $[0, \omega_{\max}]$ occupied by the dispersion therefore occupies the fraction

$$\frac{\omega_{\max}}{\omega_s} = \frac{\omega_{\max}}{\pi\,\omega_{\max}} = \frac{1}{\pi} \approx 0.3183 \tag{7}$$

of the sampling headroom: the lattice's own Nyquist frequency exceeds its propagation band by a factor $\pi \approx 3.142$, so the wave description never demands modes above the network's sampling limit. The two descriptions of the model are consistent with margin $1/\pi$.

4.3 Per-update error recursion

Let e_n be a node's estimation error just after its n-th update. Between updates the tracked instability grows at rate λ (A4), so over one interval $\Delta t = a/v$ the error is multiplied by $e^{\lambda \Delta t}$. At the update the node applies a fractional correction of gain g, replacing the error by a factor $(1-g)$ of its pre-update value. The recursion is

$$e_{n+1} = (1-g)\, e^{\lambda \Delta t}\, e_n. \tag{8}$$

Define the per-update net growth factor

$$r(g, \lambda) = (1-g)\, e^{\lambda \Delta t}. \tag{9}$$

The error diverges if $r > 1$ and decays if $r < 1$; the stability boundary is $r = 1$.

4.4 Stability threshold

Setting $r = 1$ in Eq. (9):

$$(1-g)\, e^{\lambda_{\text{crit}} \Delta t} = 1 \quad \Longrightarrow \quad \lambda_{\text{crit}} \Delta t = -\ln(1-g) \quad \Longrightarrow \quad \lambda_{\text{crit}} = -\ln(1-g)\,\frac{v}{a}. \tag{10}$$

With $v = c$ (A2) this is $\lambda_{\text{crit}} = -\ln(1-g)\, c/a$. The inputs are $c/a = 2.998 \times 10^{8} / 1.616 \times 10^{-35} = 1.855 \times 10^{43}\ \text{s}^{-1}$ and $-\ln(1-g)$. Evaluating over the gain set $g \in \{0.1, 0.5, 0.9\}$:

- $g = 0.1$: $-\ln(0.9) = 0.10536$, so $\lambda_{\text{crit}} = 0.10536 \times 1.855 \times 10^{43} = 1.954 \times 10^{42}\ \text{s}^{-1}$.
- $g = 0.5$: $-\ln(0.5) = 0.69315$, so $\lambda_{\text{crit}} = 0.69315 \times 1.855 \times 10^{43} = 1.286 \times 10^{43}\ \text{s}^{-1}$.
- $g = 0.9$: $-\ln(0.1) = 2.30259$, so $\lambda_{\text{crit}} = 2.30259 \times 1.855 \times 10^{43} = 4.271 \times 10^{43}\ \text{s}^{-1}$.

The threshold spans a factor of $2.30259/0.10536 \approx 21.9$ across this gain range, so the trackable instability rate is sensitive to the update protocol but bounded: for any $g < 1$ the threshold is finite, and finiteness of the causal speed and finiteness of the trackable instability rate stand or fall together.

4.5 Divergence timescale above threshold

Take $g = 1/2$, so $\lambda_{\text{crit}} = 1.286 \times 10^{43}\ \text{s}^{-1}$, and an instability running at $1.5\lambda_{\text{crit}} = 1.929 \times 10^{43}\ \text{s}^{-1}$. The update interval is

$$\Delta t = \frac{a}{c} = \frac{1.616 \times 10^{-35}}{2.998 \times 10^{8}} = 5.391 \times 10^{-44}\ \text{s}. \tag{11}$$

The phase accumulated per update is

$$\lambda \Delta t = 1.929 \times 10^{43} \times 5.391 \times 10^{-44} = 1.040, \tag{12}$$

so the per-update net growth factor is

$$r = (1 - 0.5)\, e^{1.040} = 0.5 \times 2.829 = 1.414. \tag{13}$$

Starting from a unit error, the number of updates needed to overrun a tolerance factor $T = 10^6$ is

$$n = \frac{\ln T}{\ln r} = \frac{\ln 10^6}{\ln 1.414} = \frac{13.8155}{0.3466} = 39.9 \approx 40 \text{ updates}. \tag{14}$$

The elapsed time is

$$t_{\text{div}} = n\,\Delta t = 39.9 \times 5.391 \times 10^{-44}\ \text{s} = 2.15 \times 10^{-42}\ \text{s}. \tag{15}$$

An instability only 50% above the threshold therefore destroys the network's estimation coherence within about 40 update intervals — a timescale of order 10^{-42} s at Planck spacing. The margin between usable and unusable tracking is narrow: at $\lambda = \lambda_{\text{crit}}$ exactly, $r = 1$ and the error neither grows nor decays, while at $1.5\lambda_{\text{crit}}$ divergence is essentially immediate on the network's own clock.

4.6 Congestion delay

Under A5 the additional propagation delay per metre is $\Delta\tau = \alpha\,\rho/\rho_c$ with $\alpha = 1/c$. At the density ratio $\rho/\rho_c = 10^{-3}$ used in [10]:

$$\Delta\tau = \frac{1}{2.998 \times 10^8\ \text{m/s}} \times 10^{-3} = 3.336 \times 10^{-12}\ \text{s/m}. \tag{16}$$

This reproduces the congestion delay derived in [10] at the same density ratio, confirming that the linear congestion convention of A5 is the one that work uses.

4.7 Device-scale cutoff

For a laboratory or device geometry, take the effective lattice spacing to be $a = 1\ \text{mm} = 10^{-3}$ m. The frequency cutoff corresponding to the zone-edge angular frequency is

$$B_{\text{mm}} = \frac{c}{2\pi a} = \frac{2.998 \times 10^8}{2\pi \times 10^{-3}}\ \text{Hz} = \frac{2.998 \times 10^8}{6.2832 \times 10^{-3}}\ \text{Hz} = 4.772 \times 10^{10}\ \text{Hz}. \tag{17}$$

This is the device-scale projection: a millimetre-spaced realization of the model would exhibit its causal ceiling in the tens-of-GHz range, which is accessible to microwave-frequency laboratory electronics. We take the prefactor $c/(2\pi a)$ under the convention that the reported cutoff is the ordinary frequency $B = \omega_{\max}/2\pi$; [10] reports 4.775×10^{10} Hz for the same millimetre-scale geometry, and the two values agree to within the rounding of the input constants.

4.8 The causal cone

Collecting the results: the lattice supports no group velocity above $v = c$ (Eq. 3), so two estimator nodes separated by $|\Delta x|$ cannot exchange any information bearing on each other's state faster than

$$|\Delta x| \le c\,\Delta t. \tag{18}$$

This is the Lorentzian causal cone in its operational form. Its isotropy — the same bound in every spatial direction — is not derived by the one-dimensional model and is carried as an assumption: a higher-dimensional lattice with isotropic low-k dispersion would be required to derive it, and the extension is listed as an open question in Section 6. What the one-dimensional derivation does establish is that the bound is exact, saturated at $k = 0$, and enforced by stability: a node that attempted to track an instability propagating faster than c/a per lattice step would need a correction gain exceeding unity, and Eq. (10) shows the error then diverges.

5. Results

The paper derives six quantitative results, each with its arithmetic shown in Section 4. They are collected here with their inputs and status.

R1 (Causal bound). For the lattice dispersion $\omega(k) = (2v/a)|\sin(ka/2)|$, the group velocity $v_g(k) = v\cos(ka/2)$ satisfies $|v_g| \le v$ for all k in the Brillouin zone, with the maximum attained at $k = 0$ (Section 4.1). This is an exact analytic result of the model, independent of the value of a.

R2 (Causal bandwidth). For $a = \ell_P = 1.616 \times 10^{-35}\ \mathrm{m}$ and $v = c = 2.998 \times 10^{8}\ \mathrm{m/s}$: $\omega_{\max} = 2c/\ell_P = 3.710 \times 10^{43}\ \mathrm{s^{-1}}$ and $B = \omega_{\max}/2\pi = 5.905 \times 10^{42}\ \mathrm{Hz}$ (Section 4.2). The result depends on a only through the scaling $\omega_{\max} \propto 1/a$.

R3 (Nyquist consistency). The network's sampling angular frequency is $\omega_s = \pi\,\omega_{\max}$, so the propagation band uses a fraction $1/\pi \approx 0.318$ of the sampling headroom (Section 4.2). The wave and estimation descriptions are consistent by a finite margin rather than marginally.

R4 (Stability threshold). The per-update recursion $e_{n+1} = (1-g)e^{\lambda\Delta t}e_n$ gives the exact threshold $\lambda_{\mathrm{crit}}(g) = -\ln(1-g)\,c/a$, evaluated as 1.954×10^{42}, 1.286×10^{43} and $4.271 \times 10^{43}\ \mathrm{s^{-1}}$ at $g = 0.1, 0.5, 0.9$ respectively (Sections 4.3–4.4). The threshold is finite for every $g < 1$ and varies by a factor of about 21.9 across the evaluated gain range.

R5 (Divergence time). At $g = 1/2$ and $\lambda = 1.5\lambda_{\mathrm{crit}}$, the per-update growth factor is $r = 1.414$, a 10^6-fold tolerance is overrun in $n = 39.9 \approx 40$ updates, and the divergence time is $t_{\mathrm{div}} = 2.15 \times 10^{-42}\ \mathrm{s}$ (Section 4.5).

R6 (Congestion and device cutoff). The congestion delay is $\Delta\tau = 3.336 \times 10^{-12}\ \mathrm{s/m}$ at $\rho/\rho_c = 10^{-3}$, reproducing [10]; the millimetre-geometry cutoff is $B_{\mathrm{mm}} = 4.772 \times 10^{10}\ \mathrm{Hz}$, consistent with the $4.775 \times 10^{10}\ \mathrm{Hz}$ reported in [10] to input rounding (Sections 4.6–4.7).

6. Discussion

Status of the conjecture. The paper establishes that a nearest-neighbor estimation lattice with finite communication speed enforces an exact causal bound, that the same lattice's update protocol enforces a finite threshold on the trackable instability rate, and that exceeding the threshold destroys estimation coherence on a timescale of tens of update intervals. What the paper does not establish is that the physical vacuum is such a lattice, that the update protocol of A4 describes any real dynamics, or that the correction gain g corresponds to a physical parameter; the same honesty boundary is drawn explicitly in [10], whose discussion states that the vacuum-lattice identification, the response-function form and the congestion-to-metric map are not established by that formalization. The present conjecture is formalized and made falsifiable; it is not confirmed.

Falsification. The main result — that c is the tracing capacity of the estimation network — would be falsified by any of the following observations. First, an analogue lattice simulation built to the specification of Section 7 in which superluminal-information propagation (communication speed above the lattice's group-velocity ceiling) does not induce estimator-error divergence: the recursion of Eq. (8) predicts divergence within $O(1/\ln r)$ updates, so a stable superluminal configuration refutes the stability mechanism. Second, a demonstration that the causal bound $|v_g| \le v$ fails in the model itself, for example under a long-range coupling that makes the dispersion $\omega(k)$ unbounded in k; the bound in Eq. (3) is exact only for the nearest-neighbor dispersion of Eq. (1). Third, a measurement of Lorentz violation with a direction-dependent causal speed, which would contradict the isotropy the model must supply in higher dimension; the extended-relativity framework of [2] shows precisely how the kinematic structure would respond to a drifting c, so a drift or anisotropy would also falsify the model's assumption of network stationarity.

Limitations. The result rests on the following assumptions, with sensitivity noted for each:

- **Lattice structure (A1).** The dispersion is specific to a nearest-neighbor scalar lattice. A different graph changes the dispersion shape and the numerical prefactor of $\omega_{\max}$, though not the existence of a ceiling, which follows from any bounded group velocity; the tree geometry of [8] is one such alternative, and that work derives its light-cone structure from Lieb–Robinson bounds on exactly such a graph. The Planck-spacing evaluation is a projection, not a derivation: the numerical values of R2, R4 and R5 scale as $1/a$ or $1/a$ with logarithmic corrections and would shift by orders of magnitude under a different effective spacing.
- **Communication speed (A2).** Identifying v with the measured c is an input, not an output; the model explains the finiteness and the bound structure, not the numerical value $2.998 \times 10^8\ \mathrm{m/s}$.
- **First-order response (A3).** The -20 dB-per-decade suppression above cutoff, the signature of early stopping identified in [10], is assumed rather than derived from the update protocol; a different transfer function changes the suppression slope but not the threshold of Eq. (10), which depends only on the recursion.
- **Update protocol (A4).** The threshold $\lambda_{\text{crit}} = -\ln(1-g)\,c/a$ is exact for fractional-correction updates with constant gain; a gain schedule, a delay, or noise in the correction changes the recursion and with it the threshold. The divergence time of R5 is sensitive to the overshoot factor: at $\lambda = 1.1\lambda_{\text{crit}}$ the phase per update is $\lambda\Delta t = 1.1 \times 0.69315 = 0.7625$, so the growth factor is $r = 0.5\,e^{0.7625} \approx 1.072$, giving $n = 13.8155/\ln(1.072) = 13.8155/0.0694 \approx 199$ updates, so the divergence time grows logarithmically as the overshoot shrinks.
- **Isotropy.** The one-dimensional model does not derive the directional universality of the cone; extending Eq. (18) to an isotropic higher-dimensional lattice is an open problem.
- **Congestion convention (A5).** The proportionality $\alpha = 1/c$ is a convention chosen to reproduce [10]; a different congestion law changes R6 but nothing upstream.

Open questions. Three questions remain open. (i) Does an isotropic higher-dimensional lattice of estimators, with the recursion of Eq. (8) at each node, reproduce the full Lorentz group rather than the single cone inequality? (ii) Does the congestion variable ρ/ρ_c map onto a geometric quantity, as the curvature-like delay of [10] suggests? (iii) What fixes the gain g physically — the answer would collapse the factor-of-22 spread of R4 to a single predicted threshold.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The primary users are researchers in discrete-substrate quantum gravity and in analogue-spacetime simulation: if the result holds, the design target of an analogue vacuum experiment shifts from reproducing kinematics to reproducing a stability budget — the experimenter would choose the lattice spacing, update rate and correction gain so that the simulated instability rates sit below $\lambda_{\text{crit}}(g) = -\ln(1-g)\,v/a$, and the experiment's causal cone would then be a consequence of the control design rather than of the chosen wave equation. A second user group is the distributed-control community: the threshold of Eq. (10) is a design inequality for any networked estimator of an unstable process, of the kind the accelerator community already engineers — the ILC final-doublet work package holds the vertical beam fluctuation at the interaction-point quadrupole stable to the order of 50 nm so that the intra-train collision feedback stays within capture range [4], and a stability-budget formalism of the form $\lambda < -\ln(1-g)/\Delta t$ makes that capture-range requirement a computable line item in the network's update protocol.

(b) Readiness. This work is purely analytical and has not been peer reviewed; under the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported, concept formulated). No claim of hardware readiness, manufacturing readiness, or experimental validation is made; the highest level this analytical work could reach is TRL 3, and only after the proposed calculation below is independently reproduced and the analogue concept is demonstrated on hardware.

(c) Smallest next test. The smallest de-risking step is a numerical calculation, not a hardware build: implement the recursion of Eq. (8) on a one-dimensional lattice of 10^3 nodes with random initial errors, drive one node with an instability at $\lambda = 0.5\lambda_{\text{crit}}$, λ_{crit} and $1.5\lambda_{\text{crit}}$ for $g = 0.5$, and measure the error norm $|e_n|$ versus update index. The prediction is that $|e_n|$ decays geometrically below threshold, stays flat at threshold, and grows by the factor $r = 1.414$ per update above threshold, reaching the 10^6 tolerance at $n \approx 40$. An observed divergence below threshold, or a bounded error above threshold, refutes the recursion and with it the stability mechanism. The next step after that is the analogue lattice simulation: a ring of physical delay-line or microwave nodes with nearest-neighbor coupling at a controlled update rate, in which the communication speed and the injected instability rate are set independently and the onset of incoherence is observed as the threshold of Eq. (10) is crossed.

(d) Binding constraints. For the analogue realization, the constraints that would bind first are: the update-rate ceiling of the physical nodes, since the threshold scales as $1/\Delta t$ and any real clock limits the trackable λ; the gain precision, since R4 spans a factor of 21.9 over $g \in [0.1, 0.9]$ and the threshold location must be resolved against it; and inter-node timing skew, which acts as an unmodeled delay in the recursion of Eq. (8) and would shift the measured threshold. Whether timing skew enters as an effective gain reduction or as a noise term is an open question the numerical test of (c) would answer first. For any claim connecting the model to physical vacuum physics, the binding constraint is conceptual: no current observation selects the lattice spacing a, and the Planck-spacing evaluation is a projection.

(e) Order-of-magnitude bounds (projections). Two projections follow from the derived scalings, with their assumptions stated. First, assuming a millimetre effective spacing (Section 4.7), the causal ceiling falls at $B_{\text{mm}} = 4.772 \times 10^{10}$ Hz, i.e. in the microwave band accessible to commodity electronics; the projection assumes the same nearest-neighbor dispersion and is uncertain by the ratio of the true effective spacing to $1\ \text{mm}$. Second, assuming the update interval equals the light-travel time across the spacing, a device with node spacing a and clock rate f satisfies the stability budget only for instabilities $\lambda < -\ln(1-g)\,2\pi f$; at $f = 10\ \text{GHz}$ and $g = 0.5$ this bound is $\lambda < 0.69315 \times 6.2832 \times 10^{10} = 4.36 \times 10^{10}\ \text{s}^{-1}$. Both projections scale linearly with their assumed inputs and carry no claim beyond the model.

(f) Risks and misuse paths. No dual-use risk is identified: the model makes no claim about physical vacuum engineering, and its control-theoretic content is a stability inequality of the standard kind already used in networked estimation. The main misuse is rhetorical — presenting the Planck-spacing numbers of R2, R4 and R5 as measured properties of nature rather than as projections of an unconfirmed model; the status statement above and the limitations list of Section 6 are the guard against that reading.

8. Conclusion

This paper has converted a conjecture — that the speed of light is the tracing capacity of a distributed-estimation vacuum — into a formal model with derived quantities and falsification conditions. The lattice estimation model yields an exact causal bound $|v_g| \le v$; a Planck-spaced realization gives a causal bandwidth $\omega_{\max} = 3.710 \times 10^{43}\ \text{s}^{-1}$ and $B = 5.905 \times 10^{42}$ Hz with Nyquist margin $1/\pi$; the per-update error recursion gives the exact stability threshold $\lambda_{\text{crit}}(g) = -\ln(1-g)\,c/a$, equal to $1.286 \times 10^{43}\ \text{s}^{-1}$ at $g = 1/2$ and spanning a factor of 21.9 over $g \in [0.1, 0.9]$; an instability at $1.5\lambda_{\text{crit}}$ overruns a 10^6-fold tolerance in 2.15×10^{-42} s, about 40 update intervals; and the congestion delay $3.336 \times 10^{-12}\ \text{s/m}$ at $\rho/\rho_c = 10^{-3}$ and the millimetre-geometry cutoff 4.772×10^{10} Hz reproduce the corresponding results of [10]. The conjecture is formalized and made falsifiable; it is not confirmed. The numerical test specified in Section 7 — iterating the error recursion across the threshold — is the smallest step that would confirm or refute the stability mechanism, and the analogue lattice simulation built on it would be the first physical setting in which a causal cone is produced by a control design rather than postulated as kinematics.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.2). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 15 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Topological Control of Chirality and Spin with Structured Light. arXiv:2508.08733v3. https://arxiv.org/abs/2508.08733v3 [2] Special and general relativity extended to allow for a variable light speed. arXiv:astro-ph/0606542v4. https://arxiv.org/abs/astro-ph/0606542v4 [3] On The Varying Speed of Light in a Brane-Induced FRW Universe. arXiv:hep-th/9912037v2. https://arxiv.org/abs/hep-th/9912037v2 [4] Midterm Status Report of the ILC Technology Network Activities. arXiv:2603.01172v1. https://arxiv.org/abs/2603.01172v1 [5] StegNet: Mega Image Steganography Capacity with Deep Convolutional Network. arXiv:1806.06357v1. https://arxiv.org/abs/1806.06357v1 [6] Light-shift spectroscopy of optically trapped atomic ensembles. arXiv:1911.02708v1. https://arxiv.org/abs/1911.02708v1 [7] Spatial, spectral, temporal and polarisation resolved state tomography of light. arXiv:2202.01932v1. https://arxiv.org/abs/2202.01932v1 [8] https://papers.qnfo.org/papers/formal-derivation-of-emergent-spacetime. QNFO: Formal Derivation of Emergent Spacetime to the Bruhat-Tits Tree. [9] https://papers.qnfo.org/papers/syntactic-generation-primitive-distinctions. QNFO: Syntactic Generation. [10] https://papers.qnfo.org/papers/the-speed-of-light-as-a-regularization-parameter-of-a-distributed-estimation-vac. QNFO: The Speed of Light as a Regularization Parameter of a Distributed-Estimation Vacuum: Causal Bandwidth, Early-Stopping Suppression, and Latency Congestion. [11] https://papers.qnfo.org/papers/all-things-are-one. QNFO: All Things Are One.