

The Speed of Light as a Regularization Parameter of a Distributed-Estimation Vacuum: Causal Bandwidth, Early-Stopping Suppression, and Latency Congestion

The Speed of Light as a Regularization Parameter of a Distributed-Estimation Vacuum: Causal Bandwidth, Early-Stopping Suppression, and Latency Congestion

Abstract

The speed of light c is measured, codified in the SI, and used daily, but it is normally treated as a kinematic axiom rather than as a parameter with a physical function. This paper develops and quantifies the conjecture that a finite causal speed acts as a regularization parameter of a vacuum modeled as a distributed estimation network, suppressing high-frequency divergences the way early stopping suppresses overfitting in iterative learning systems. We formalize the vacuum as a lattice of locally coupled estimators with signal velocity v , derive the dispersion relation $\omega(k) = (2v/a) \sin(ka/2)$, and show that the causal bound $|v_g| \leq v$ yields a maximum angular frequency $\omega_{\max} = 2v/a = 3.71 \times 10^{43} \text{ s}^{-1}$ and a channel bandwidth $B = 5.905 \times 10^{42} \text{ Hz}$ for a lattice spacing at the Planck length. A spectral response model $S(\omega) = (1 + (\omega/\omega_c)^2)^{-1}$ suppresses modes at $10\omega_c$ by a factor of 9.90×10^{-3} and modes at $100\omega_c$ by 9.999×10^{-5} , a roll-off of approximately -20 dB per decade. A congestion model in which latency scales as $\tau = (L/v)(1 + \rho/\rho_c)$ reproduces curvature-like delay effects parametrically, with a worked example giving $\Delta\tau = 3.336 \times 10^{-12} \text{ s}$ per metre at $\rho/\rho_c = 10^{-3}$. The result reframes c as a stability parameter and specifies simulations and laboratory tests that would confirm or refute the conjecture.

1. Introduction

Physicists, metrologists, and quantum-engineering researchers face a persistent explanatory gap: the speed of light $c = 2.998 \times 10^8 \text{ m/s}$ is measured and used daily, but no existing framework explains why the causal speed is finite at all, or what physical role finiteness plays beyond kinematics. This gap has practical consequences. In quantum error correction, engineering teams must allocate physical-to-logical qubit ratios of roughly 1000:1 to fight decoherence that is treated as an irreducible feature of continuous-state quantum mechanics [8]; if fragility is instead a symptom of an under-regularized substrate, the design target for error-correction overheads changes. In cosmology, theorists must accept that the cosmic microwave background is isotropic to a few parts in 10^5 across regions causally disconnected at recombination [3], and resolve this horizon problem by adding an inflationary epoch; a framework in which the causal speed itself is a regulated, environment-dependent parameter offers an alternative lever, and variable-speed-of-light constructions have already been embedded consistently in extended special and general relativity [2] and in string-theoretic brane dynamics [3].

What decision would change if the conjecture of this paper holds? If a finite signal velocity provably suppresses high-frequency instabilities in a distributed estimation network — the same role early stopping plays in iterative learning systems — then c acquires a functional interpretation: it is the regularization parameter that keeps the vacuum's estimation dynamics from diverging on high-frequency fluctuations, allowing only low-frequency, stable structures to propagate as observable reality. Gravitational effects become reinterpretable, at least parametrically, as latency congestion in a bandwidth-limited causal channel. This gives theorists a concrete stability criterion — a

critical bandwidth window — to test in lattice simulations, and gives experimentalists dispersion signatures to search for, of the kind already predicted in pre-geometric frameworks where high-energy photons accumulate propagation delays of $\Delta t = (E/E_P) \times 10^{-4}$ seconds per kiloparsec [11].

The idea has precedent in structure, though it is new in function. Discrete-geometry work has derived an invariant light speed from the spectral radius of the Bruhat-Tits tree, identifying the dimensionless upper bound of c with a Lieb-Robinson velocity and concluding that c is a derived, macroscopic, epistemic property of combinatorial geometry rather than a fundamental axiom [9]. Those works derive the value of c ; none of them assigns c a regulatory role in stabilizing the substrate's own dynamics. That assignment is the contribution here.

The contributions of this paper are:

- A formal lattice estimation model of the vacuum in which a finite signal velocity v acts as an early-stopping constraint, with the dispersion relation $\omega(k) = (2v/a) \sin(ka/2)$ and the causal bound $|v_g| \leq v$ derived explicitly (Sections 3 and 4.1).
- A numerical derivation of the causal bandwidth: $\omega_{\max} = 2v/a = 3.71 \times 10^{43} \text{ s}^{-1}$ and $B = 5.905 \times 10^{42} \text{ Hz}$ for a at the Planck length and $v = c$ (Section 4.2).
- A spectral suppression analysis showing that a bandwidth-limited response $S(\omega)$ attenuates modes at $10\omega_c$ by 9.90×10^{-3} and at $100\omega_c$ by 9.999×10^{-5} , establishing the early-stopping analogy quantitatively (Section 4.3).
- A stability criterion — a critical bandwidth window inside which low-frequency structures survive and outside which modes are suppressed — stated as an explicit inequality on growth versus damping rates (Section 4.4).
- A congestion model $\tau = (L/v)(1 + \rho/\rho_c)$ that reproduces curvature-like delay effects parametrically, with a worked example at $\rho/\rho_c = 10^{-3}$ (Section 4.5), plus a device-scale variant of the cutoff applied to cold-atom trap geometries (Section 4.7).
- A simulation protocol, a laboratory falsification protocol, and engineering pathways for the conjecture (Sections 5, 6, and 7).

2. Background and Related Work

This section reviews the works that constrain what a "speed of light as regularization parameter" model may claim. Three threads are relevant: variable and emergent light speed, discrete and combinatorial substrates, and analogies from estimation theory and optical physics.

Variable light speed as a consistent extension of relativity. The variable-speed-of-light (VSL) literature establishes that c can be treated as a dynamical or environmental parameter without abandoning relativity. [2] extends the Lorentz transformation, the Minkowski line element, and the field equation of general relativity to allow a variable light speed $c(t)$ in a homogeneous and isotropic universe, under the minimal assumption that all moving observers at the same spacetime point measure the same light speed even though it may vary; the construction introduces the differential $d\hat{T} = c dT$ and a rest mass m that varies so as to keep the rest energy $\hat{m} = mc^2$ constant, yielding covariant velocity, energy-momentum, and stress-energy tensors. This matters here because it removes the objection that treating c as a parameter is automatically inconsistent: the kinematic machinery survives a variable c . Our model goes further in interpretation — c is not merely allowed to vary but is assigned a stabilizing function — and the framework of [2] is the natural consistency check for any variation we posit.

VSL realized in string theory. [3] realizes the varying-speed-of-light scenario concretely: a 3+1 dimensional probe brane in the background of a black hole formed by a stack of branes, analyzed via the Dirac-Born-Infeld action, in which the speed of light on the probe brane decreases in time as the brane moves toward the bulk black hole. The

authors add a rotating black-hole background to prevent the brane from hitting the horizon, bulk expansion as effective friction, and freezeout with virialization to stabilize the orbit, ensuring the speed of light tends to a constant at late times. Two features of this work feed directly into our model. First, it demonstrates a physical mechanism by which the local causal speed is set by the substrate's state — exactly the dependence our congestion model requires, where v_{eff} depends on a local congestion density. Second, the same work frames the horizon problem quantitatively: the cosmic microwave background is isotropic to a few parts in 10^5 across regions causally disconnected at recombination, which standard causal microphysics cannot explain. In our reading, a bandwidth-limited estimation substrate with a globally shared regularization parameter is a candidate mechanism for correlated behavior across causal distances, because the constraint is a property of the network, not of signal exchange between regions.

Emergent light speed from discrete geometry. [9] derives macroscopic spacetime and the invariant speed of light from the discrete Bruhat-Tits tree: the tree's adjacency operator acts as an effective Hamiltonian in a rigged Hilbert space, time emerges as an epistemic coarse-grained parameter via the Page-Wootters mechanism and a modular Hamiltonian, and the dimensionless upper bound of the speed of light is the Lieb-Robinson velocity derived from the graph's spectral radius, with Lorentz invariance emerging as a continuum symmetry of the Archimedean projection. The stated conclusion is that " c is a derived, macroscopic, epistemic property of combinatorial geometry." Our work is complementary rather than competing: [9] derives the value and invariance of c from graph structure; we ask what that finite value does — and answer that it bounds the substrate's estimation bandwidth, which is precisely the role a Lieb-Robinson-type velocity plays in limiting how fast correlations can spread on a graph. The Lieb-Robinson bound in [9] is, in our language, the early-stopping constraint made geometric.

Fragility of the continuous substrate and the signal-processing ceiling. [8] motivates discrete substrates from an engineering failure mode: quantum states decohere rapidly not merely from technical noise but from a mathematical vulnerability of the continuous real-number field underlying standard quantum mechanics, and active quantum error correction imposes unsustainable burdens — surface codes requiring physical-to-logical ratios of roughly 1000:1, thermodynamic walls from Landauer's principle and Heisenberg back-action, and a signal-processing ceiling beyond which no classical filtering removes fundamental quantum back-action. This supports our premise that the vacuum behaves like an estimation system under noise whose continuous limit is unstable without regularization. [8] also lists testable signatures of a discrete substrate — deviations at ultra-low energies, anisotropies in the fine-structure constant, discrete-spacetime effects in cosmic rays, and modified dispersion relations at high energies — and modified dispersion is exactly the observable class our bandwidth window predicts (Section 4.3).

Constants as geometric scaling factors, with falsifiable dispersion. [11] proposes that reality emerges from a pre-geometric S^1 substrate with integer winding numbers, reinterpreting fundamental constants as geometric scaling factors, and makes falsifiable propagation predictions: a modulation in inspiraling binary gravitational waves at $f_{\text{mod}} = \frac{m_1 m_2}{m_1 + m_2} \times f_{\text{orbital}}$, and an energy-dependent photon dispersion $\Delta t = (E/E_P) \times 10^{-4}$ seconds per kiloparsec, testable against gamma-ray propagation and the gravitational-wave detection record. This demonstrates the methodological standard our conjecture must meet: a reinterpretation of constants is publishable physics only when it emits a dispersion signature. Section 4.3 derives the shape of such a signature for the regularization model.

Estimation-theoretic analogy: truncated horizons and measurability failures. The dynamic-programming literature supplies the formal analogy for early stopping. [5] proves that the value function of a discounted dynamic program satisfies the Bellman equation under stated regularity conditions on the primitives, and documents the failure modes that motivate regularization: even with Borel-measurable primitives, the value function need not be Borel measurable, with a concrete example in which $S = A = [0, 1]$, a Borel set $B \subseteq S \times A$ whose projection E is not Borel, and reward $r = 1_B(a, s)$, so that $Tf(s) = \sup_{a \in A} 1_B(a, s) = 1_E(s)$ is not Borel measurable. The lesson we import is that unregularized optimization over a rich state space can produce ill-defined

objects; finite-horizon truncation is a standard repair. A vacuum estimating an unstable stochastic process faces the same pathology, and a finite propagation speed is a physical implementation of horizon truncation: no estimator can condition on information farther than vT away, so the effective optimization horizon is finite and set by c .

Optical analogies: propagation-induced structure, bandwidth-limited encoding, compensation, and characterization. Four optics papers supply mechanistic analogies rather than direct support. [1] shows that spin-orbit interaction and optical chirality can arise during propagation itself, exclusively from the Pancharatnam topological index ℓ_p , with no pre-imposed modal asymmetry at the generation plane — the input beam carries no initial spin and identical radial amplitudes for both circular components, yet deterministic spin and chirality emerge from propagation. This is a concrete optical instance of the shape of our central claim: a propagation constraint (there topological, here causal) deterministically generates structure that was not present in the initial state. [4] demonstrates that a spatially varying ac-Stark shift maps each atom's radial position to a unique frequency shift, so that a single light-shift-broadened spectrum encodes the spatial distribution of the ensemble, with the trap itself providing the reference. This is the encoding principle behind our congestion model: position-dependent shifts of a shared reference encode where in a potential field an element sits, which is how a latency map could encode a gravitational potential. [7] provides the compensation dual: a weak additional field tuned to suppress the influence of a strong trapping field narrows a strongly distorted optical transition, with results over a range of compensation detunings matching a numerically evaluated density-matrix model. In our framework, this is the experimental pattern for testing whether adding a controlled bandwidth constraint suppresses an unwanted spectral distortion — the direct laboratory analogue of regularization. Finally, [6] demonstrates full state tomography of light — spatial, spectral, temporal, and polarization — using SLM-encoded projective measurements and intensity-only detection through single-mode fibre, capable of analyzing mutually incoherent fields within one beam. This is the measurement toolchain a laboratory test of our dispersion predictions would use: any bandwidth-induced modification of a probe beam's state is in principle fully characterizable by these methods.

A caution on numerology. [10] argues that base-10's dominance in science stems from human anatomy rather than mathematical naturalness, and that this anthropocentric contingency can masquerade as fundamental structure. We take this as a methodological warning for the present program: claims that constants "encode" computationally meaningful values must be checked against the possibility that the pattern is an artifact of the representational system used to describe them. Our derivations therefore use dimensionless ratios wherever possible.

3. Methods

3.1 Model definition

We model the vacuum as a regular d -dimensional lattice of locally coupled estimators. Each site i at lattice spacing a holds a local estimate $x_i(t)$ of an underlying stochastic field and exchanges information only with nearest neighbors, with a signal velocity v per link. The model's assumptions are stated explicitly:

- **A1 (discreteness).** The substrate is a lattice with spacing a . We take a to be the Planck length $\ell_P = 1.616 \times 10^{-35}$ m as a reference value; the derivation holds for any a .
- **A2 (finite causal speed).** Information propagates at finite speed v ; no estimator receives information from beyond radius vT in time T . This is the early-stopping constraint under test.
- **A3 (local estimation dynamics).** Each site updates its estimate from neighbor information and local noise; the collective dynamics admit wave-like excitations with dispersion $\omega(k)$.
- **A4 (unstable target process).** The underlying stochastic process has unstable high-frequency components, i.e., modes whose naive amplitude would grow without the causal constraint.

Assumption A2 is the physical content of the conjecture; A1 and A3 are the standard apparatus of lattice field models; A4 encodes the "overfitting" pathology the regularization is meant to prevent.

3.2 Dispersion and the causal bound

For a nearest-neighbor coupled lattice, the standard dispersion relation for scalar excitations is

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

where k is the wavenumber restricted to the first Brillouin zone $|k| \leq \pi/a$. The group velocity is

$$v_g(k) = \frac{\partial \omega}{\partial k} = v \cos\left(\frac{ka}{2}\right),$$

so that $|v_g(k)| \leq v$ for all k , with equality only at $k = 0$. This is the lattice realization of a Lieb-Robinson-type bound of the kind [9] identifies as the origin of the invariant speed: the maximum correlation-spreading velocity is a property of the graph and its coupling, not of the excitations.

3.3 Spectral response and suppression model

To quantify the early-stopping analogy we posit a bandwidth-limited response function for the substrate, of the standard first-order low-pass form,

$$S(\omega) = \frac{1}{1 + (\omega/\omega_c)^2},$$

where ω_c is the cutoff angular frequency of the causal channel. This form is an assumption (a modeling choice, not derived from A1–A4); its role is to convert the causal bound into a concrete suppression factor per frequency decade, which is what a simulation would measure.

3.4 Stability criterion

Let a mode of angular frequency ω have an intrinsic (unconstrained) amplitude growth rate $\gamma(\omega) \geq 0$ and let the causal constraint impose an effective damping $\Gamma(\omega)$. A mode is stable within the substrate if

$$\Gamma(\omega) > \gamma(\omega).$$

We adopt the minimal polynomial forms $\gamma(\omega) = \gamma_0(\omega/\omega_c)^2$ and $\Gamma(\omega) = \Gamma_0$, both assumptions, giving the stability window

$$\omega < \omega_c \sqrt{\Gamma_0/\gamma_0} \equiv \omega_s.$$

Only modes below ω_s propagate as stable, observable structure; modes above are suppressed — the "only low-frequency stable signals propagate as reality" claim of the originating conjecture, now an inequality.

3.5 Congestion model

We model gravitational effects as latency congestion. Let ρ be a local density of active estimation load and ρ_c a saturation capacity, both dimensionless after normalization. The effective signal velocity is

$$v_{\text{eff}} = \frac{v}{1 + \rho/\rho_c},$$

and the latency across a path of length L is

$$\tau(\rho) = \frac{L}{v_{\text{eff}}} = \frac{L}{v} \left(1 + \frac{\rho}{\rho_c} \right).$$

The fractional excess latency $\Delta\tau/\tau_0 = \rho/\rho_c$ is the model's curvature-like observable: regions of high estimation load behave as regions of increased optical depth. This is a parametric claim only; mapping ρ/ρ_c onto the metric tensor of general relativity is future work (Section 6).

3.6 Simulation protocol

The conjecture is tested by varying v in lattice or agent-based simulations of the model of Section 3.1 and measuring: (i) the amplitude of high-frequency modes as a function of v ; (ii) the survival of localized, low-frequency structures; (iii) delay accumulation in inhomogeneous load distributions. The predictions and falsification conditions are in Sections 5 and 6.

4. Analysis

All input numbers are stated with their origin. Derived quantities are computed step by step.

4.1 The causal bound as an early-stopping constraint

From the dispersion relation of Section 3.2,

$$v_g(k) = v \cos\left(\frac{ka}{2}\right),$$

the extrema follow from $\cos(ka/2) = \pm 1$, i.e., $k = 0$ and $k = \pm 2\pi/a$ (outside the first Brillouin zone for the latter). Within $|k| \leq \pi/a$, the maximum of $|v_g|$ is at $k = 0$:

$$|v_g|_{\text{max}} = v \cos(0) = v.$$

Hence the causal bound $|v_g| \leq v$ holds for every mode. Interpretation: no excitation, however energetic, can propagate information faster than the substrate's coupling velocity. This is the physical early-stopping rule — it truncates the effective estimation horizon of every site to vT , exactly as finite-horizon truncation bounds the path length explored in dynamic programming, the setting in which [5] shows unregularized value functions can fail to be Borel measurable.

4.2 Causal bandwidth of the vacuum channel

Inputs. Lattice spacing $a = \ell_P = 1.616 \times 10^{-35}$ m (assumption A1, reference value; the Planck length is used as the conventional scale at which continuum spacetime is expected to fail, and no result below depends on this choice beyond rescaling). Signal velocity $v = c = 2.998 \times 10^8$ m/s (the measured value; the SI value is exact by definition, and we use three significant figures).

Step 1: maximum angular frequency. The dispersion $\omega(k) = (2v/a)|\sin(ka/2)|$ is maximized at the Brillouin-zone edge $k = \pi/a$, where $\sin(\pi/2) = 1$:

$$\omega_{\text{max}} = \frac{2v}{a} = \frac{2 \times 2.998 \times 10^8 \text{ m/s}}{1.616 \times 10^{-35} \text{ m}}.$$

Numerator: $2 \times 2.998 \times 10^8 = 5.996 \times 10^8$ m/s. Division:

$$\omega_{\max} = \frac{5.996 \times 10^8}{1.616 \times 10^{-35}} = 3.710 \times 10^{43} \text{ s}^{-1}.$$

Step 2: channel bandwidth in hertz. The ordinary-frequency bandwidth is

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

Step 3: interpretation. A causal channel with speed c on a Planck-spaced lattice can carry at most $B = 5.905 \times 10^{42} \text{ Hz}$ of mode bandwidth. Any estimation dynamics requiring bandwidth above this is not representable by the substrate — the finite- c constraint is a hard spectral ceiling, not a soft preference. This is the quantitative form of "the vacuum cannot overfit high-frequency fluctuations": it literally lacks the bandwidth to represent them as propagating degrees of freedom.

4.3 Suppression factors of the bandwidth-limited response

With the response model $S(\omega) = (1 + (\omega/\omega_c)^2)^{-1}$ of Section 3.3, we compute the suppression at integer multiples of the cutoff. We identify ω_c with $\omega_{\max} = 3.710 \times 10^{43} \text{ s}^{-1}$ from Step 1 (the most conservative choice: the cutoff is the physical ceiling itself).

At $\omega = 10 \omega_c$:

$$S(10 \omega_c) = \frac{1}{1 + 10^2} = \frac{1}{101} = 9.901 \times 10^{-3}.$$

At $\omega = 100 \omega_c$:

$$S(100 \omega_c) = \frac{1}{1 + 100^2} = \frac{1}{10001} = 9.999 \times 10^{-5}.$$

Suppression ratio per decade: the response falls by a factor

$$\frac{S(10 \omega_c)}{S(100 \omega_c)} = \frac{9.901 \times 10^{-3}}{9.999 \times 10^{-5}} = 99.02 \approx 10^2$$

per decade of frequency above cutoff — i.e., approximately -20 dB per decade, the canonical signature of a first-order regularization filter. A simulation varying v would measure precisely this roll-off in the mode-amplitude spectrum; observing a roll-off shallower than first-order would falsify the specific response model while leaving the bandwidth-ceiling claim of Section 4.2 intact.

4.4 The critical bandwidth window

With $\gamma(\omega) = \gamma_0(\omega/\omega_c)^2$ and $\Gamma(\omega) = \Gamma_0$ (Section 3.4), the stability boundary $\Gamma_0 = \gamma_0(\omega_s/\omega_c)^2$ gives

$$\omega_s = \omega_c \sqrt{\frac{\Gamma_0}{\gamma_0}}.$$

As a worked reference point, take the balanced case $\Gamma_0 = \gamma_0$ (assumption; the physical case of interest is where the substrate's damping and the process's instability are of comparable scale):

$$\omega_s = \omega_c \sqrt{1} = \omega_c = 3.710 \times 10^{43} \text{ s}^{-1}.$$

The stable window is then $0 \leq \omega < \omega_c$: all modes below the causal ceiling propagate stably, all modes at or above it are suppressed. For a general ratio $r = \Gamma_0/\gamma_0$, the window edge scales as $\sqrt{r}$: a substrate with weaker damping ($r = 0.01$) stabilizes only modes below $0.1 \omega_c$, while stronger damping ($r = 100$) stabilizes up to $10 \omega_c$ — but the latter is unphysical here because the channel itself carries no content above $\omega_{\max} = \omega_c$. The critical bandwidth window is therefore generically of order ω_c , set by the causal speed, and its width in ordinary frequency is $B = 5.905 \times 10^{42}$ Hz (Section 4.2). This is the paper's central stability claim: stable macroscopic structure requires modes inside a window whose upper edge is fixed by v , so varying v in a simulation should move the edge and change which structures survive.

4.5 Congestion and curvature-like delay

Worked example. Take a path of length $L = 1$ m, baseline velocity $v = c = 2.998 \times 10^8$ m/s, and a local congestion fraction $\rho/\rho_c = 10^{-3}$ (assumption for illustration).

Baseline latency:

$$\tau_0 = \frac{L}{v} = \frac{1}{2.998 \times 10^8} = 3.336 \times 10^{-9} \text{ s.}$$

Congested latency:

$$\tau = \tau_0 \left(1 + \frac{\rho}{\rho_c}\right) = 3.336 \times 10^{-9} \times 1.001 = 3.339 \times 10^{-9} \text{ s.}$$

Excess delay:

$$\Delta\tau = \tau - \tau_0 = 3.336 \times 10^{-9} \times 10^{-3} = 3.336 \times 10^{-12} \text{ s.}$$

The effective refractive index of the congested region is

$$n_{\text{eff}} = \frac{c}{v_{\text{eff}}} = 1 + \frac{\rho}{\rho_c} = 1.001.$$

A spatial gradient in ρ/ρ_c therefore produces a position-dependent delay of exactly the kind that [4] exploits experimentally, where a position-dependent light shift maps radial position to a unique frequency offset and the trap itself provides the reference. In our model, the congestion field plays the role of the trap potential: a latency map imprinted on the causal channel. Whether this map satisfies an equation of the Einstein type is not derived here; what is derived is that the model produces curvature-like observables (position-dependent delays and effective indices) with a single new parameter ρ/ρ_c .

4.6 Consistency with the variable- c formalism

The congestion model makes v_{eff} a function of the substrate state, which is a variable light speed in the sense formalized by [2]: the extended Lorentz transform and Minkowski line element there accommodate any c under the assumption that all observers at the same spacetime point measure the same light speed. Our $v_{\text{eff}}(\rho)$ satisfies that assumption locally. Likewise, the brane realization of [3], in which the probe-brane light speed decreases as the brane approaches a bulk black hole and tends to a constant at late times through friction and virialization, is a physical precedent for a causal speed that is set by the substrate's configuration and that stabilizes to a fixed value — the behavior our congestion model exhibits when the estimation load saturates.

4.7 Device-scale variant of the cutoff (labeled projection)

The Planck-scale result of Section 4.2 fixes the fundamental ceiling. A second, device-scale application of the same framework uses a macroscopic characteristic interaction length L in place of the lattice spacing. We assume a first-order response with cutoff frequency

$$f_c = \frac{c}{2\pi L},$$

analogous to the -20 dB-per-decade roll-off of Section 4.3 but with the cutoff set by the device geometry. This is a projection under the stated assumption, not a consequence of the Planck-scale derivation.

Inputs. $c = 3.00 \times 10^8$ m/s (three significant figures); $L = 1.0 \times 10^{-3}$ m (assumed transverse radius of a dipole trap, an order-of-magnitude scale consistent with the cold-atom ensembles studied in [4]); $h = 6.626 \times 10^{-34}$ J s (physical constant).

Step 1: cutoff frequency.

$$f_c = \frac{c}{2\pi L} = \frac{3.00 \times 10^8}{2\pi \times 1.0 \times 10^{-3}} = \frac{3.00 \times 10^8}{6.2832 \times 10^{-3}} = 4.775 \times 10^{10} \text{ Hz},$$

$$f_c = \frac{3.00 \times 10^8}{6.2832 \times 10^{-3}} = 4.775 \times 10^{10} \text{ Hz}.$$

Step 2: equivalent energy scale. With $h = 6.626 \times 10^{-34}$ J s,

$$E_c = hf_c = 6.626 \times 10^{-34} \times 4.775 \times 10^{10} = 3.164 \times 10^{-23} \text{ J}.$$

Converting to electron-volts with $1 \text{ eV} = 1.602 \times 10^{-19} \text{ J}$,

$$E_c = \frac{3.164 \times 10^{-23}}{1.602 \times 10^{-19}} = 1.975 \times 10^{-4} \text{ eV}.$$

Step 3: interpretation. A millimetre-scale device therefore realizes a regularization cutoff at $f_c \approx 4.775 \times 10^{10}$ Hz, i.e., at microwave-to-millimetre-wave frequencies and an energy scale of 2×10^{-4} eV, which is in the range accessible to cold-atom hyperfine and microwave spectroscopy of the kind performed in dipole-trapped ensembles [4]. This device-scale projection is what makes the conjecture laboratory-testable at all: the Planck-scale ceiling of Section 4.2 is 10^{33} times higher in frequency, but the same geometric scaling $f \sim c/(2\pi L)$ maps the cutoff onto any chosen length scale. We stress that this scaling is an assumption of the device variant, adopted for testability; the fundamental derivation fixes only the Planck-scale value.

5. Results

The paper derives five 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 $|v_g| \leq v$ for all $|k| \leq \pi/a$, 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 $a = \ell_P = 1.616 \times 10^{-35}$ m and $v = c = 2.998 \times 10^8$ m/s:

$$\omega_{\max} = \frac{2c}{\ell_P} = 3.710 \times 10^{43} \text{ s}^{-1}, \quad B = \frac{\omega_{\max}}{2\pi} = 5.905 \times 10^{42} \text{ Hz}.$$

The inputs are the reference lattice spacing (assumption A1) and the measured light speed; the arithmetic is in Section 4.2. The result depends on a only through the scaling $\omega_{\max} \propto 1/a$.

R3 (Suppression roll-off). With the assumed first-order response $S(\omega) = (1 + (\omega/\omega_c)^2)^{-1}$ and $\omega_c = \omega_{\max}$: $S(10\omega_c) = 1/101 = 9.901 \times 10^{-3}$, $S(100\omega_c) = 1/10001 = 9.999 \times 10^{-5}$, and the decade ratio $S(10\omega_c)/S(100\omega_c) = 99.02 \approx 10^2$, i.e., approximately -20 dB per decade (Section 4.3). The response shape is a modeling assumption; the roll-off exponent is what a simulation of Section 3.6 would measure first.

R4 (Stability window). With $\gamma(\omega) = \gamma_0(\omega/\omega_c)^2$ and $\Gamma(\omega) = \Gamma_0$, the stable window is $\omega < \omega_s = \omega_c \sqrt{\Gamma_0/\gamma_0}$; in the balanced case $\Gamma_0 = \gamma_0$ the window edge equals the causal ceiling $\omega_c = 3.710 \times 10^{43} \text{ s}^{-1}$, and for a general ratio $r = \Gamma_0/\gamma_0$ the edge scales as $\sqrt{r} \omega_c$ (Section 4.4).

R5 (Congestion delay). For $L = 1 \text{ m}$, $v = c$, and $\rho/\rho_c = 10^{-3}$: $\tau_0 = 3.336 \times 10^{-9} \text{ s}$, $\tau = 3.339 \times 10^{-9} \text{ s}$, $\Delta\tau = 3.336 \times 10^{-12} \text{ s}$ per metre, and $n_{\text{eff}} = 1.001$ (Section 4.5). The congestion fraction is an assumed illustration; the derived content is the linear relation $\Delta\tau/\tau_0 = \rho/\rho_c$.

R6 (Device-scale cutoff, labeled projection). Under the assumed scaling $f_c = c/(2\pi L)$ with $L = 1.0 \times 10^{-3} \text{ m}$: $f_c = 4.775 \times 10^{10} \text{ Hz}$ and $E_c = hf_c = 3.164 \times 10^{-23} \text{ J} = 1.975 \times 10^{-4} \text{ eV}$ (Section 4.7). This is a projection under the stated scaling assumption, not a consequence of R2.

No empirical measurement is reported in this paper. Every number above is either an exact consequence of the stated model equations (R1, R4), an arithmetic consequence of stated inputs (R2, R3, R5), or a projection under a labeled assumption (R6).

6. Discussion

Status of the conjecture. The paper establishes that a finite causal speed, embedded in a nearest-neighbor estimation lattice, acts as a hard spectral ceiling and, under an assumed first-order response, as a -20 dB-per-decade suppressor of above-cutoff modes — the quantitative signature of early stopping. What the paper does not establish is that the physical vacuum is such a lattice, that the response function has the assumed form, or that the congestion variable ρ/ρ_c maps onto the metric tensor. The conjecture is formalized and made falsifiable; it is not confirmed.

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

- **Lattice structure (A1).** The dispersion $\omega(k) = (2v/a)|\sin(ka/2)|$ is specific to a nearest-neighbor scalar lattice. A different graph (e.g., a different coordination number or a tree geometry as in [9]) 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 bandwidth value $B = 5.905 \times 10^{42} \text{ Hz}$ is sensitive to this choice at the order-of-magnitude level.
- **Planck spacing.** Identifying a with ℓ_P is a convention. All bandwidth numbers rescale as ℓ_P/a ; the paper's claims about suppression and stability are independent of this choice.
- **Response function.** $S(\omega) = (1 + (\omega/\omega_c)^2)^{-1}$ is posited, not derived. A second-order response would give -40 dB per decade; a sharp brick-wall cutoff would give infinite attenuation above ω_c . The early-stopping analogy is quantitative only under the first-order choice.
- **Growth and damping forms.** $\gamma \propto \omega^2$ and constant Γ are minimal assumptions. The window edge $\omega_s = \omega_c \sqrt{\Gamma_0/\gamma_0}$ moves as $\sqrt{r}$, so a ratio r differing by a factor of 10^4 moves the edge by 10^2 ; the qualitative claim (a finite window) survives, the quantitative edge does not.
- **Congestion linearity.** $v_{\text{eff}} = v/(1 + \rho/\rho_c)$ is the simplest saturating form. Any monotone decreasing $v_{\text{eff}}(\rho)$ yields curvature-like delays; only the linear fractional excess $\Delta\tau/\tau_0 = \rho/\rho_c$ is specific to this

choice.

- **No field-equation map.** The model produces position-dependent delays and effective indices, not Einstein's equations. Deriving $G_{\mu\nu} = 8\pi GT_{\mu\nu}$ -type dynamics from a congestion field is open.

Failure modes. Three are foreseeable. First, a lattice estimation model may develop artifacts — numerical dispersion, aliasing at the Brillouin-zone edge — that mimic suppression without representing physics; simulations must distinguish model artifacts from the claimed regularization. Second, the congestion model may be unidentifiable: many latency fields produce the same delay map, so a measured $\Delta\tau$ field does not uniquely determine ρ . Third, the analogy to early stopping is structural, not a theorem: no result here proves a generalization-error bound for the vacuum estimation network of the kind proven for iterative learning on statistical models.

What would falsify the main result. The central claim is that varying the signal velocity v in the simulation protocol of Section 3.6 moves the stability edge and changes which structures survive, with a mode-amplitude roll-off of approximately first order above the cutoff. The single observation that would falsify the specific mechanism is a simulation in which increasing v does not shift the survival boundary of high-frequency structures, or in which the measured roll-off is flat (no suppression) above the causal ceiling while stable structures still fail to form. On the observational side, the model shares falsification conditions with the broader discrete-substrate program: [8] lists modified dispersion relations at high energies and Lorentz-violation signatures among the testable predictions of a discrete substrate, and [11] predicts an energy-dependent photon delay $\Delta t = (E/E_P) \times 10^{-4}$ seconds per kiloparsec; a precision test finding no dispersion at sensitivities exceeding these predictions would not refute the regularization conjecture in its generality, but would eliminate the specific dispersion signatures this class of models currently offers.

Open questions. (i) Does a rigorous early-stopping bound exist for distributed estimation on a graph with bounded information velocity, analogous to the guarantees known for iterative learning? (ii) Can the congestion field be coupled to stress-energy so that the weak-field limit reproduces Newtonian gravity with the measured G ? (iii) Does the horizon problem admit a bandwidth-based resolution — correlated initial conditions as a property of a shared, globally regularized substrate — that survives quantitative comparison with the few-parts-in- 10^5 isotropy constraint [3]? (iv) What fixes Γ_0/γ_0 physically?

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The primary users at this stage are theorists working on quantum gravity, discrete-substrate models, and the foundations of quantum mechanics, and simulation groups running lattice field or agent-based models. If the conjecture holds in simulation, the design target changes in two places: lattice studies of emergent spacetime gain a stability criterion (the critical bandwidth window) to impose on candidate substrates, and the interpretation of c in variational frameworks — already treated as a dynamical parameter in the consistent extensions of [2] and realized microphysically in the brane construction of [3] — acquires a functional justification for why the parameter is finite. A secondary, longer-horizon user community is quantum-error-correction engineering: if decoherence fragility is a symptom of an under-regularized continuous substrate, as the resource analysis of [8] argues (surface-code physical-to-logical ratios of roughly 1000:1, thermodynamic walls, a signal-processing ceiling), then a substrate-level regularization would change the overhead budget, though no such engineering consequence follows from the present analysis and none is claimed.

(b) Current readiness. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 — basic principles observed and reported, with the model equations formulated and the key quantities derived on paper, but no simulation execution, no laboratory proof of concept, and

no review. Analytical work of this kind cannot exceed TRL 3, and the device-scale projection of Section 4.7 is itself a paper construct, not a built apparatus.

(c) The smallest next test. The cheapest de-risking step is a 1D or 2D numerical lattice simulation of the model of Section 3.1: evolve coupled estimator sites with signal velocity v as a control parameter, inject a known unstable high-frequency mode spectrum, and measure the mode-amplitude roll-off above the cutoff $\omega_{\max} = 2v/a$. The falsifying outcome is stated in Section 6: no shift of the survival boundary as v varies, or a flat (unattenuated) spectrum above the ceiling. The smallest laboratory follow-on, using the device-scale projection R6, is a cold-atom experiment in the dipole-trap geometry of the kind characterized in [4] and compensated in [7]: probe the ensemble with a controlled bandwidth-limited field and test whether imposing the bandwidth constraint suppresses an unwanted spectral distortion in the manner the compensation technique of [7] already demonstrates for light shifts. Full characterization of any modified probe state would use the spatial, spectral, temporal, and polarization tomography demonstrated in [6].

(d) Binding constraints. For the simulation pathway, the binding constraint is scale separation: demonstrating that the stability edge tracks v requires resolving modes across many decades below $\omega_{\max}$, which stresses time-stepping and grid resolution. For the laboratory pathway, the binding constraints are the achievable bandwidth and stability of probe fields at the 10^{10} – 10^{11} Hz scale of R6, and the identification problem noted in Section 6: a measured delay map does not uniquely determine the congestion field. Whether the congestion parameter can be coupled to real stress-energy is an open question, not a derived constraint.

(e) Order-of-magnitude bounds (projections). Under the stated assumptions only: the Planck-scale channel bandwidth is $B = 5.905 \times 10^{42}$ Hz (derived, R2); the device-scale cutoff for $L = 1$ mm is $f_c = 4.775 \times 10^{10}$ Hz with equivalent energy 1.975×10^{-4} eV (projection, R6, assuming $f_c = c/(2\pi L)$); the per-metre congestion delay at $\rho/\rho_c = 10^{-3}$ is 3.336×10^{-12} s (worked illustration, R5). No cost figures are derivable from the present analysis and none are asserted.

(f) Risks and misuse paths. The main misuse is interpretive overreach: presenting the congestion model as a theory of gravity or the bandwidth ceiling as an established property of spacetime, when the paper derives only parametric analogies under stated assumptions. A second risk is numerological overfitting of the kind [10] warns against — reading fundamental significance into representational coincidences of the constants; the derivations here use dimensionless ratios and stated assumptions specifically to limit this, but any future mapping of ρ/ρ_c onto gravitational parameters must clear the same bar.

8. Conclusion

This paper has converted a conjecture — that the speed of light functions as a regularization parameter of a distributed-estimation vacuum — into a formal model with derived quantities and falsification conditions. The lattice estimation model yields an exact causal bound $|v_g| \leq v$; 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; an assumed first-order response gives a -20 dB-per-decade suppression of above-cutoff modes, the quantitative signature of early stopping; a stability window $\omega_s = \omega_c \sqrt{\Gamma_0/\gamma_0}$ states the survival criterion as an inequality; and a congestion model produces curvature-like delays, $\Delta\tau = 3.336 \times 10^{-12}$ s per metre at $\rho/\rho_c = 10^{-3}$, with a device-scale projection placing a laboratory-accessible cutoff at 4.775×10^{10} Hz (1.975×10^{-4} eV) for millimetre-scale geometry. The framework is consistent with the variable- c formalisms of [2] and [3], complements the derivation of c from discrete geometry in [9] by assigning the derived finite speed a stabilizing function, and inherits testable dispersion

signatures from the discrete-substrate program of [8] and [11]. The conjecture remains unconfirmed; its value is that it is now stated precisely enough to be refuted — by a simulation in which varying v fails to move the stability edge, or by dispersion searches that find no effect at the predicted parametric level.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.18.6-codeagent). 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, zai-org/glm-5.3, openai/gpt-oss-120b, 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: a separately generated script recomputed 11 quantitative claims from the paper's stated inputs; 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] Light-shift spectroscopy of optically trapped atomic ensembles. arXiv:1911.02708v1. <https://arxiv.org/abs/1911.02708v1> [5] The Principle of Optimality in Dynamic Programming: A Pedagogical Note. arXiv:2302.08467v5. <https://arxiv.org/abs/2302.08467v5> [6] Spatial, spectral, temporal and polarisation resolved state tomography of light. arXiv:2202.01932v1. <https://arxiv.org/abs/2202.01932v1> [7] Dual-colour magic-wavelength trap for suppression of light shifts in atoms. arXiv:1811.04508v1. <https://arxiv.org/abs/1811.04508v1> [8] QNFO: Alpha Pi Project [9] QNFO: Formal Derivation of Emergent Spacetime to the Bruhat-Tits Tree [10] QNFO: Ten-Fingered Trap [11] QNFO: Winding Number as Hidden Variable