

Network Tracing Capacity as an Informational Event Horizon: A Delay-Based Causal Bound for Networked Control

Network Tracing Capacity as an Informational Event Horizon: A Delay-Based Causal Bound for Networked Control

Abstract

Distributed control systems—power-electronics clusters, vehicle platoons, coordinated robot fleets—stabilize unstable plants through communication networks, and the latency of that communication, not the control law, frequently decides whether stabilization succeeds. Existing design practice treats this as a per-application tuning problem; no topology-level invariant is known that states, in one number, which instability rates a given communication topology can and cannot propagate. We propose such an invariant, the Network Tracing Capacity (NTC), defined for a benchmark class of scalar unstable modes under delayed feedback as $\Theta = 1/(h\tau)$, where h is the controller-to-sensor hop count and τ the per-hop delay, and we conjecture that modes with growth rate $\lambda > \Theta$ are practically unobservable through the topology, a boundary we interpret as an informational analog of a relativistic event horizon. We build the correspondence mapping (instability rate, delay, hop count) onto (expansion rate, signal speed, metric scale), anchor its operational criteria in how the Event Horizon Telescope actually detects and is limited by horizons, and derive a worked example: a five-hop network with $\tau = 10\text{ ms}$ yields $\Theta = 20\text{ s}^{-1}$, so a plant with $\lambda = 35\text{ s}^{-1}$ lies beyond the horizon at hop 3 of 5, with state growth factor $e^{1.75} \approx 5.75$ per delay interval. Reducing the per-hop delay to 2 ms restores $\Theta = 100\text{ s}^{-1} > \lambda$. A complementary topology-attribute convention, based on boundary-crossing-path walk counts, is derived and compared. The result, if validated, converts latency budgets into hop-count-aware instability-rate budgets for control topology design.

1. Introduction

Engineers who close control loops over communication networks face a failure mode that no amount of controller gain can remove. When the plant is unstable with growth rate λ and the measurement of its state must travel h communication hops, each with delay τ , the controller acts on information about a state that has already grown by a factor $e^{\lambda h\tau}$ since it was measured. Beyond a critical product of instability rate, hop count, and per-hop delay, the loop diverges regardless of the algorithm at either end. The practitioners who hit this wall are the designers of networked control systems: distributed power-electronics clusters, platooning vehicles, teleoperated and swarm robotics, and any plant whose sensors and actuators are separated by a multi-hop network. Today the boundary is found per deployment, empirically, by tuning and testing; a design that works on one topology silently fails when a routing change adds a hop. What is missing is a single, topology-level number that tells the designer, before deployment, which instability rates the topology can propagate and which it structurally cannot.

If such a number exists and has the character we conjecture—an informational event horizon, analogous to the causal horizon that the speed of light c creates in spacetime—then three things change. First, controllability and observability analysis acquires a topology-independent bound: the question "can this plant be stabilized over this network?" becomes a comparison of two scalars, λ and Θ , computable from the graph and its latencies alone. Second, the design lever shifts from controller synthesis to topology and latency engineering, because Θ depends

only on h and τ . Third, control-theoretic observability and relativistic causality would be unified under one invariance principle: in both settings, a finite propagation speed creates a boundary beyond which information about a region ceases to be recoverable from observations.

This paper is a first, deliberately modest step toward that conjecture. We do not prove a general theorem for all networked control systems; we construct the benchmark class in which the claim is derivable, derive the threshold with explicit arithmetic, and state precisely what would falsify the analogy. Our contributions are:

- We define the Network Tracing Capacity $\Theta(G, \tau) = 1/(h_{\max}\tau)$ for a communication topology G with uniform per-hop delay τ , and the horizon condition $\lambda < \Theta$, under a stated benchmark criterion for delayed stabilization of a scalar unstable mode (Section 3, formalized and derived in Section 4).
- We construct the explicit correspondence mapping from (instability rate λ , per-hop delay τ , hop count h) onto (expansion rate, invariant signal speed, metric scale), and identify the operational criteria any horizon analog must satisfy (Section 3).
- We derive, with full arithmetic, a worked example: for $h = 5$ hops and $\tau = 10$ ms, $\Theta = 20$ s⁻¹; a plant with $\lambda = 35$ s⁻¹ crosses the horizon at hop 3, and the state grows by $e^{1.75} \approx 5.75$ across the full delay; reducing τ to 2 ms restores $\Theta = 100$ s⁻¹ $> \lambda$ (Sections 4 and 5).
- We derive a complementary topology-attribute capacity $\kappa_{\text{NTC}}^{(k)}$ from the boundary-crossing-path emergence discrepancy of [10], obtaining $\kappa_{\text{NTC}}^{(1)} = 4$ for single-step observations, and document the convention choice between the two definitions (Sections 3.6, 4, and 6).
- We assess the analogy against the operational practice of horizon science—how the Event Horizon Telescope detects horizon-scale structure and where its inference demonstrably degrades—and state the falsification condition for the control-side claim (Sections 2, 6, and 7).

2. Background and Related Work

The physical side of our analogy rests on how event horizons are actually observed, not merely defined. The Event Horizon Telescope (EHT) is a global very long baseline interferometry (VLBI) array operating at 1.3 mm wavelength, composed of millimeter/submillimeter telescopes scattered across the globe—from Hawaii and Arizona to Chile, Mexico, the South Pole, and Spain and France—with the primary goal of imaging supermassive black holes at a few gravitational radii r_g [1]. Two features of this program matter for us. First, the horizon manifests observationally as a stable geometric scale: the 1.3 mm emission regions in Sgr A and M87 have an effective size of only a few r_g , and their small sizes combined with weak infrared emission are the core of the observational case that these sources possess event horizons [1]. Second, the horizon-scale structure is measured to be stable across years: reanalyses of M87 across the 2017, 2018, and 2021 epochs find ring diameters of $[42.0, 46.4]$ μas in 2017, $[40.7, 44.4]$ μas in 2018, and $[43.1, 44.4]$ μas in 2021, consistent across methods and epochs, exactly as accretion theory predicts for a geometry-set scale [5]. A horizon, operationally, is a boundary whose geometric signature is invariant even as the plasma near it varies—a property our proposed Θ threshold is meant to share.

The same literature also demonstrates where inference near a horizon degrades, which is the behavior our unobservability claim mirrors. The first EHT analysis of circular polarization in M87 developed five independent imaging approaches, all validated on synthetic data from GRMHD simulations, and all found a moderate degree of image circular polarization; yet the reconstructed circular polarization structures vary considerably among the five methods, and the collaboration concludes that the structure cannot be reliably inferred without additional data or stronger assumptions [2]. For Sgr A, dynamic imaging and geometric snapshot modeling of minute-timescale variability recover the position angle in some synthetic GRMHD movies but produce prominently biased results in others [3]. This is precisely the phenomenology of an informational boundary: observables that straddle the horizon

carry information that reconstruction cannot fully recover, and different pipelines disagree because the missing information is not a noise problem but a causal one. The Sgr A campaign itself illustrates the cost of fighting variability across a delay-limited aperture: Sgr A is highly variable, with flickering and flares on short timescales, and among known black holes it has the largest predicted angular size, approximately $50 \mu\text{as}$ [6]; the 2017 imaging relied on the best-coverage day, April 7, because April 11 showed the highest light-curve variability, rendering static imaging for that day particularly challenging [3]. Calibration of such variable sources required novel methods and two independent pipelines, EHT-HOPS and rPICARD [6].

Horizon-scale imaging also feeds back into fundamental tests, which shows what a mature "horizon metrology" can deliver. The Sgr A images constrain deviations from the Kerr metric—the only stationary, vacuum, axisymmetric, uncharged solution of Einstein's equations without pathologies outside the event horizon—and the combination of imaging with infrared limits strengthens the argument that Sgr A lacks a thermalizing surface, because the discrepancy between the maximum steady infrared luminosity the source can have and the minimum luminosity a surface would produce is too large to be circumvented by small model fixes [8]. The mid-range science program extends this metrology toward dynamics, including the untested prediction that the M87* jet is powered by the black hole's spin through an analog of the Penrose process, while the operational cycle—observations proposed a year ahead, nearly a year to assemble recorded data at correlator sites, and another year for calibration and validation—shows how severely finite information transport rates bind even flagship science [4]. Polarized emission at 230 GHz around M87, imaged on event-horizon scales, probes magnetic-field structure and plasma properties near the horizon [7], completing the picture of a boundary that is probed, bounded, and exploited rather than merely postulated.

On the network side, the closest precedent for cross-domain structural claims is the framework of network isomorphism, which defines isomorphism as exact structural equivalence—two networks are isomorphic if one transforms into the other by moving and relabeling nodes without altering connections—and argues that equivalence between complex systems cannot be captured by a single measure but must be assessed progressively across layers of topological structure, functional architecture, and dynamic behavior, packaged methodologically as the L1–L2–L3 isomorphism hierarchy and the Abstract Network Object [9]. Our correspondence mapping is a dynamic-layer (L3) claim in exactly that sense: we do not claim that a communication topology is geometrically isomorphic to spacetime, but that a dynamically defined quantity—the propagatable instability rate—maps onto a dynamically defined relativistic quantity. A second precedent establishes that physics–network-engineering isomorphisms can be made productive: the topological extension framework models entities as finite-volume stable modes of nonlinear evolution equations, asserts that "existence" is functionally equivalent to "persistence" of a mode against dispersive forces, and maps the screening mechanisms of cosmological chameleon fields onto route flap damping protocols in interdomain routing, yielding a structural isomorphism in which stability is maintained by a conserved topological charge and environmental coupling [11]. Our horizon conjecture follows the same pattern—identify a mechanism (causal suppression of instability) that recurs in two domains and give it one name.

Finally, the methodological template for turning a structural analogy into checkable arithmetic comes from the boundary-crossing-path theory of network emergence. That work proves, verified by exhaustive enumeration on a reference network, that for route-accumulating observations the emergence discrepancy equals the boundary-touching walk count exactly, with per-length values $D^{(1)} = 4$, $D^{(2)} = 12$, $D^{(3)} = 28$ (cumulative 44) and closed-form growth $D_{\leq L} = 2^{L+3} - 8 - 4L$, giving $D_{\leq 5} = 220$ and $D_{\leq 10} = 8144$; its discussion is explicit that the identity holds for a narrow class of observations and that weighted or stochastic interactions would require replacing walk counts by weighted path sums [10]. We adopt the same discipline: derive the bound exactly in a narrow benchmark class, state the class boundary, and treat generalization as an open question rather than an assertion.

What no cited work provides—and what motivates this paper—is a quantitative, topology-level threshold at which a networked control loop transitions from observable to unobservable, phrased so that the threshold can be compared, term by term, with a relativistic horizon. The EHT literature supplies the operational anatomy of a horizon; the network-isomorphism literature supplies the licensing for cross-domain mappings; the missing piece is the control-theoretic derivation, which Sections 3 and 4 supply for the benchmark class.

3. Methods

3.1 Benchmark plant and communication model

We consider the minimal object to which the conjecture can be applied without ambiguity: a scalar unstable plant mode

$$\dot{x}(t) = \lambda x(t) + u(t), \quad \lambda > 0,$$

whose state $x(t)$ is measured at a sensor node and must reach a controller node across a communication topology G along a path of h hops, each hop imposing a deterministic transport delay τ . The total information delay is

$$\tau_d = h\tau.$$

We assume deterministic delays, no packet loss, and a single dominant unstable mode; these restrictions are load-bearing and are revisited in Section 6.

3.2 The benchmark stabilization criterion

The derivation requires a criterion that separates "stabilizable over the network" from "not stabilizable over the network." We adopt the following as the defining criterion of the benchmark class, stated openly as a modeling assumption rather than as established literature:

Assumption A1 (delay-budget criterion). A delayed-feedback scheme for the scalar mode maintains bounded estimation and control error only if the uncontrolled growth accumulated over one full information delay is at most a factor e ; that is,

$$\lambda\tau_d \leq 1.$$

The rationale is a growth-versus-correction budget: over one delay interval τ_d , the uncontrolled state grows by $e^{\lambda\tau_d}$, and the controller must undo this accumulated growth using information that is itself τ_d old. We take unity exponent as the boundary of the benchmark class because it is the value at which the growth over one delay equals the natural exponential scale of the mode itself; the qualitative structure of everything below is unchanged if the constant is replaced by any fixed κ , with the threshold rescaled to $\kappa/(h\tau)$. Choosing the constant is a modeling decision; the paper's claim is about the structure of the threshold, not the value of its constant.

3.3 Network Tracing Capacity and the horizon condition

Under Assumption A1, the largest instability rate the topology can propagate is

$$\Theta(G, \tau) = \frac{1}{h\tau} = \frac{1}{\tau_d},$$

which we call the Network Tracing Capacity (NTC) of the topology. The horizon condition is

$$\lambda < \Theta \iff \lambda h\tau < 1.$$

Modes with $\lambda \geq \Theta$ are declared beyond the horizon in the benchmark class: no controller at the far end of an h -hop, τ -delay path can stabilize them, because the information on which it acts describes a state that has grown by more than the criterion allows.

3.4 The correspondence mapping

The conjectured isomorphism assigns, term by term:

Networked control	Relativistic spacetime
Instability rate λ	Local expansion rate of the congruence of worldlines
Per-hop delay τ , hop count h	Metric scale; invariant speed c sets τ per unit distance
NTC $\Theta = 1/(h\tau)$	Invariant speed c as the propagatable rate bound
Hop index k with $\lambda k \tau = 1$	Horizon radius: the last surface sending signals that arrive
$\lambda \tau_d > 1$: unobservable mode	Region beyond the horizon: causally disconnected

The mapping is structural, not physical: no claim is made that network delay is Lorentz-invariant or that instability is geometric expansion. The claim is that both systems instantiate the same mechanism—a finite propagation rate interacting with a super-linear-growing quantity creates a sharp boundary beyond which information about the growing quantity cannot reach the observer in usable form.

3.5 Operational criteria for a horizon analog

Borrowing the operational anatomy of horizon science (Section 2), we require that a genuine informational horizon satisfy three criteria: (O1) the boundary is a function of structure and propagation rate alone, not of the observer's algorithm; (O2) the boundary's location is stable under variation of conditions away from the boundary, as the M87 *ring diameter is stable across 2017–2021* [5]; (O3) *inference about states beyond the boundary degrades in a way that is pipeline-dependent rather than noise-limited, as the M87 circular polarization structure "cannot be reliably inferred without additional data or stronger assumptions" across five independent methods* [2]. Criterion O1 is built into the definition of Θ ; O2 and O3 are what the simulation program of Section 7 must test.

3.6 A complementary topology-attribute convention

An alternative definition assigns the NTC not from hop count and delay but from a graph attribute: the boundary-crossing-path emergence discrepancy $D^{(k)}$ of [10], which for route-accumulating observations equals the number of boundary-touching walks of length k . Under that convention the horizon inequality reads $\lambda \tau \leq \kappa_{\text{NTC}}^{(k)}$ with $\kappa_{\text{NTC}}^{(k)} = D^{(k)}$. The two conventions differ in what plays the role of the horizon: hop-delay budget (Section 3.3) versus walk-count capacity across a cut. Section 4 derives the walk-count values, and the choice of the delay-budget convention for the main text is documented on its merits in Sections 3.6 and 6.

4. Analysis

Every number in this section is derived from the inputs stated here: $h = 5$ hops, per-hop delay $\tau = 10 \text{ ms} = 0.01 \text{ s}$ (Scenario 1), $\tau = 2 \text{ ms} = 0.002 \text{ s}$ (Scenario 2), and plant instability rate $\lambda = 35 \text{ s}^{-1}$. The criterion constant is 1 (Assumption A1).

Step 1: Total information delay, Scenario 1.

$$\tau_d = h\tau = 5 \times 0.01 \text{ s} = 0.05 \text{ s} = 50 \text{ ms}.$$

Step 2: Network Tracing Capacity, Scenario 1.

$$\Theta = \frac{1}{\tau_d} = \frac{1}{0.05 \text{ s}} = 20 \text{ s}^{-1}.$$

Step 3: Horizon test. The plant has $\lambda = 35 \text{ s}^{-1}$. Compare:

$$\lambda\tau_d = 35 \times 0.05 = 1.75 > 1,$$

so $\lambda = 35 \text{ s}^{-1} > \Theta = 20 \text{ s}^{-1}$: the mode lies beyond the horizon in the benchmark class.

Step 4: Growth factor across the full delay. The uncontrolled state grows, over one full information delay, by

$$e^{\lambda\tau_d} = e^{1.75}.$$

Computing: $e^{1.75} = e^1 \cdot e^{0.75}$; $e^{0.75} \approx 2.117000$ (since $e^{0.7} \approx 2.013753$ and $e^{0.05} \approx 1.051271$, so $e^{0.75} \approx 2.013753 \times 1.051271 \approx 2.117000$). Hence

$$e^{1.75} \approx 2.718282 \times 2.117000 \approx 5.754603 \approx 5.75.$$

Each delay interval, the controller's information describes a state **5.75** times smaller than the current one, and the deficit compounds every interval—this is the divergence mechanism behind the horizon.

Step 5: Location of the horizon in hop space. The horizon condition in hop index k is $\lambda k\tau < 1$, i.e., $k < 1/(\lambda\tau)$. Compute:

$$\frac{1}{\lambda\tau} = \frac{1}{35 \times 0.01} = \frac{1}{0.35} = 2.857142 \dots \approx 2.86.$$

Hops $k = 1$ and $k = 2$ give $\lambda k\tau = 0.35$ and 0.70 , both below 1: information arriving from one or two hops away is inside the horizon. Hop $k = 3$ gives $\lambda k\tau = 1.05 > 1$: the horizon sits between hop 2 and hop 3, so the effective horizon hop is

$$h^* = \left\lceil \frac{1}{\lambda\tau} \right\rceil = 3.$$

Of the $h = 5$ hops, $h - h^* = 2$ hops lie beyond the horizon: state information originating at those distances cannot be used for stabilization in the benchmark class, even though the packets physically arrive.

Step 6: The topology lever, Scenario 2. Reduce the per-hop delay to $\tau = 2 \text{ ms} = 0.002 \text{ s}$ at unchanged hop count:

$$\tau_d = 5 \times 0.002 = 0.01 \text{ s}, \quad \Theta = \frac{1}{0.01} = 100 \text{ s}^{-1}.$$

Now $\lambda = 35 \text{ s}^{-1} < \Theta = 100 \text{ s}^{-1}$, and the horizon test gives $\lambda\tau_d = 35 \times 0.01 = 0.35 < 1$: the same plant, same controller, same hop count, is inside the horizon. The growth factor per delay drops from $e^{1.75} \approx 5.75$ to

$$e^{0.35} \approx 1.419068 \approx 1.42.$$

Step 7: Scaling law. From $\Theta = 1/(h\tau)$, the NTC scales inversely in each factor: doubling the hop count halves Θ ; halving the per-hop delay doubles it. Formally, for a fractional change,

$$\frac{\Delta\Theta}{\Theta} \approx -\frac{\Delta h}{h} - \frac{\Delta\tau}{\tau},$$

to first order. The horizon hop moves as $h^* = \lceil 1/(\lambda\tau) \rceil$, independent of total hop count: adding hops never rescues a mode once the per-distance budget $\lambda\tau \geq 1$ is exceeded at the required distance—it only pushes more of the path beyond the horizon.

Step 8: Walk-count capacity from the boundary-crossing-path reference network. The reference network of [10] yields, for route-accumulating observations, the per-length emergence discrepancies

$$D^{(1)} = 4, \quad D^{(2)} = 12, \quad D^{(3)} = 28,$$

with cumulative $D^{(1)} + D^{(2)} + D^{(3)} = 4 + 12 + 28 = 44$, and the closed-form cumulative growth

$$D_{\leq L} = 2^{L+3} - 8 - 4L.$$

Evaluating at $L = 5$:

$$D_{\leq 5} = 2^8 - 8 - 20 = 256 - 28 = 228.$$

The closed form evaluated at $L = 10$ gives

$$D_{\leq 10} = 2^{13} - 8 - 40 = 8192 - 48 = 8144,$$

which matches the value 8144 reported in [10]. At $L = 5$, [10] reports $D_{\leq 5} = 220$ for the reference network's two-edge boundary, while the closed form as stated gives 228; we adopt the reported value 220 as the topology attribute and record the 8-unit difference as an open item on the exact boundary accounting, discussed further in Section 6. Under the walk-count convention of Section 3.6, the single-step capacity is therefore

$$\kappa_{\text{NTC}}^{(1)} = D^{(1)} = 4,$$

and the five-step cumulative capacity is $\kappa_{\text{NTC}}^{(5)} = 220$. Under that convention and a link delay $\tau = 1 \text{ ms} = 10^{-3} \text{ s}$, the corresponding rate-scale values are

$$\frac{\kappa_{\text{NTC}}^{(1)}}{\tau} = \frac{4}{10^{-3}} = 4 \times 10^3 \text{ s}^{-1}, \quad \frac{\kappa_{\text{NTC}}^{(5)}}{\tau} = \frac{220}{10^{-3}} = 2.2 \times 10^5 \text{ s}^{-1}.$$

These are attributes of a different convention, not of the delay-budget Θ ; they are reported for comparison and are not used in the main horizon test.

Step 9: What the derivation establishes and what it does not. Steps 1–7 establish: within the benchmark class (scalar mode, deterministic delays, Assumption A1), the topology admits a single propagatable-rate number Θ , a sharp horizon condition, a computable horizon hop h^* , and a topology-only lever. Step 8 establishes that a graph-attribute capacity can be computed exactly for a reference network from published walk-count results. The derivations do not establish that Assumption A1 holds for general plants, vector states, stochastic delays, or packet loss, nor that the walk-count and delay-budget conventions coincide; those extensions are open (Section 6). The strong conjecture of the research program—that the boundary is a structural analog of a relativistic event horizon rather than an engineering limitation—is supported here only to the extent that the benchmark threshold exhibits criteria O1–O3; O1 holds by construction, O2 and O3 are untested and are the target of the falsification test in Section 7.

5. Results

All numbers below are computed in Section 4 from the stated inputs; none are measured or simulated.

R1 (NTC of the reference topology). For $h = 5$ hops and $\tau = 10$ ms, the Network Tracing Capacity is $\Theta = 20 \text{ s}^{-1}$ (Section 4, Step 2).

R2 (Horizon crossing). A plant mode with $\lambda = 35 \text{ s}^{-1}$ exceeds the NTC ($\lambda\tau_d = 1.75 > 1$) and is beyond the horizon in the benchmark class (Step 3).

R3 (Growth factor). Across one full information delay, the uncontrolled state grows by $e^{1.75} \approx 5.75$ (Step 4).

R4 (Horizon hop). The horizon sits between hops 2 and 3: $h^* = \lceil 1/(0.35) \rceil = 3$, so 2 of the 5 hops lie beyond the horizon (Step 5).

R5 (Topology lever). At $\tau = 2$ ms and unchanged $h = 5$, $\Theta = 100 \text{ s}^{-1} > \lambda = 35 \text{ s}^{-1}$ and the growth factor per delay falls to $e^{0.35} \approx 1.42$: the same plant is inside the horizon (Step 6).

R6 (Scaling law). $\Theta = 1/(h\tau)$ scales inversely in hop count and per-hop delay to first order (Step 7).

6. Discussion

The central claim of this paper is deliberately narrow: within a benchmark class (a scalar unstable mode, deterministic per-hop delays, no packet loss, and the delay-budget criterion of Assumption A1), a communication topology admits a single topology-level number $\Theta = 1/(h\tau)$ that separates propagatable from non-propagatable instability rates, and this threshold exhibits the operational signature of a horizon analog. The claim is not that all networked control systems obey this bound, nor that the analogy to relativistic causality is more than structural. This section states what would falsify the claim, where the analogy is weakest, and which assumptions carry the result.

Falsification condition. The main result would be falsified by a demonstration, experimental or analytical, that a scalar unstable mode with $\lambda \geq \Theta = 1/(h\tau)$ is stabilized over an h -hop, τ -delay network by a controller that uses only the delayed measurements, under the benchmark assumptions (deterministic delays, no loss, bounded control effort). Equivalently, if a simulation in the style proposed in Section 7 shows estimation error remaining bounded at $\lambda\tau_d$ substantially greater than 1 — say $\lambda\tau_d = 3$ or more, well beyond the boundary — then Assumption A1 does not describe a real transition and the horizon reading collapses into an arbitrary normalization. Conversely, if bounded-error operation is observed at $\lambda\tau_d = 0.9$ but fails at $\lambda\tau_d = 1.1$ across independent implementations, the sharp-threshold structure survives and the research program advances. A second, softer falsifier concerns the analogy itself: if the transition at $\lambda\tau_d = 1$ turns out to be gradual (error growth accelerating smoothly rather than changing character), the threshold remains useful as an engineering budget but loses the horizon interpretation, because a genuine horizon analog should show the pipeline-dependent, not noise-limited, degradation of inference that the M87* circular polarization analysis exhibits — five independent methods all detecting a signal yet producing reconstructions that "vary considerably" and cannot be reliably inferred without stronger assumptions [2].

Assumption A1 is the load-bearing modeling choice. The criterion $\lambda\tau_d \leq 1$ was adopted because it makes the growth over one delay interval equal to the natural exponential scale of the mode; any fixed constant κ would rescale the threshold to $\kappa/(h\tau)$ without changing the structure. This means the *numerical* value of Θ is convention-dependent, and only the existence and monotonicity of the threshold are claims of the paper. A reviewer who prefers $\kappa = 2$ or $\kappa = 1/2$ obtains a different boundary; the paper's contribution is the form $\Theta \propto 1/(h\tau)$, not the constant. This must be stated plainly because it limits what the worked example shows: the example demonstrates the arithmetic of the threshold, not a measured transition.

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

- *Scalar, single-mode plant.* Real plants are vector-valued with multiple modes; the bound is applied per mode, and mode coupling could allow a super-threshold mode to be stabilized through interaction with a sub-threshold one. Sensitivity: unknown; no coupling analysis is attempted here.
- *Deterministic, uniform per-hop delays.* Real networks have jitter and queueing delay; if τ is random with mean $\bar{\tau}$, the effective budget involves a tail quantity, not the mean. Sensitivity: high, because the threshold is a sharp inequality and jitter smears it.
- *No packet loss.* Loss effectively increases the information delay by the retransmission interval; the bound extends naturally by replacing τ with an effective per-hop delay, but this extension is asserted, not derived.
- *The delay-budget criterion (A1).* As above, the constant is conventional; the structure is not.
- *Static topology.* The bound assumes the routing path is fixed; a network that reroutes mid-operation changes h and therefore Θ discontinuously, which is exactly the failure mode motivating the work but is not modeled.
- *The walk-count comparison (Step 8).* The closed form as written, $D_{\leq L} = 2^{L+3} - 8 - 4L$, evaluates at $L = 5$ to $2^8 - 8 - 20 = 228$, while [10] reports $D_{\leq 5} = 220$ for the reference network's two-edge boundary; the 8-unit difference is an unresolved discrepancy in the boundary-edge accounting convention that we have not resolved. We adopt the reported value 220 as the topology attribute and flag the inconsistency, consistent with Step 8 of Section 4.

Where the analogy is weakest. In general relativity, the horizon is observer-independent and encoded in the geometry; here, Θ depends on the chosen control path, and a different path on the same graph has a different Θ . The mapping in Section 3.4 assigns hop count to metric scale, but a network offers many paths where spacetime offers one. The strongest defense is that the *per-distance* budget $\lambda\tau$ is path-independent: the horizon hop $h^* = \lceil 1/(\lambda\tau) \rceil$ depends only on the local product, matching the local character of a relativistic horizon, while the global quantity Θ is path-averaged. Whether the per-distance budget is the deeper invariant is an open question. A second weakness: the EHT evidence for horizon-like behavior includes geometric stability across epochs — ring diameters consistent within a few microarcseconds across 2017, 2018, and 2021 [5] — and our Θ has no comparable stability test yet, since no simulation or measurement has been run. Criteria O2 and O3 of Section 3.5 remain untested.

Open questions. (i) Does the threshold persist for vector plants with structured actuators? (ii) What replaces $\lambda\tau_d \leq 1$ under stochastic delay — is the correct budget a tail probability $\Pr[\tau_d > 1/\lambda] \leq \epsilon$? (iii) Do the delay-budget and walk-count conventions of Section 3.6 coincide on any nontrivial class of graphs, or are they genuinely different quantities that merely share a name? (iv) Does the claimed connection to regularization/early-stopping rules in RKHS learning, noted in the originating research program, survive formalization? None of these is resolved here.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users, if the result is validated, are designers of networked control systems: engineers specifying communication architectures for distributed power-electronics clusters, vehicle platoons, and teleoperated robot fleets. The decision that changes is the latency budget itself. Today a latency budget is stated as a maximum end-to-end delay; under the horizon condition it becomes a joint budget on the product $h\tau$: a designer who cannot reduce τ further can instead reduce h (fewer routing hops between sensor and controller), and the bound $\lambda < 1/(h\tau)$ states exactly how much hop reduction buys how much instability-

rate headroom. A routing change that adds a hop is no longer a silent risk but a computable one: the new Θ follows from the new h in one division. This is a design-rule change, not a tooling change, and it applies at the architecture-review stage before any controller is synthesized.

(b) Current readiness. This work is purely analytical and has not been peer reviewed; against the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles observed and reported, with the concept and its applications formulated but no proof-of-concept demonstrated. The derivation of Θ and the horizon condition is analytical; analytical work never exceeds TRL 3, and this work has not yet reached TRL 3 because no simulation or experiment has tested whether the predicted transition occurs. No manufacturing readiness is claimed or implied.

(c) The smallest next experiment. The minimal de-risking test is a numerical simulation, requiring no hardware: integrate the scalar plant $\dot{x} = \lambda x + u$ with a delayed-feedback controller (for example, a delayed observer-based law) over a range of $\lambda\tau_d$ from 0.5 to 3, with h and τ varied independently to confirm the product dependence, and measure whether the bounded-error/unbounded-error transition of the estimation error occurs near $\lambda\tau_d = 1$. The outcome that refutes the claim is a transition location that depends on h and τ separately rather than on their product, or a transition that is absent (bounded error at $\lambda\tau_d = 3$). A second, equally cheap test is robustness to jitter: replace the deterministic τ with a random delay of the same mean and check whether the transition smears or shifts.

(d) Binding constraints. In order of likely first contact: (i) *delay variability* — the bound assumes deterministic τ , and real networks' queueing jitter is the first assumption to break in any deployment; (ii) *packet loss* — retransmissions inflate the effective information delay and are unmodeled; (iii) *topology dynamics* — rerouting changes h mid-operation, and the bound gives no guidance on transient behavior during the change; (iv) *the criterion constant* — until the simulation program of (c) runs, the numerical budget $\lambda\tau_d \leq 1$ is a convention, and safety-critical deployments would need a measured, derated constant. Whether mode coupling in vector plants can rescue super-threshold modes is an open question, not a quantified constraint.

(e) Order-of-magnitude bounds. As a labeled projection with stated assumptions: if the per-hop delay achievable on a dedicated industrial network is $\tau = 100 \mu\text{s} = 10^{-4} \text{s}$ (an assumption about deployment hardware, not derived here) and the plant requires $\lambda = 100 \text{s}^{-1}$, then the maximum tolerable hop count is $h < 1/(\lambda\tau) = 1/(10^{-2}) = 100$ hops — comfortably permissive. At $\tau = 10 \text{ms}$ (the Scenario 1 value derived in Section 4), the same plant tolerates only $h < 1/(0.35) \approx 2.86$, i.e., $h \leq 2$ hops. The projection spans two orders of magnitude in permissiveness across plausible per-hop delays, which is the practical content of the scaling law R6: the binding constraint in any application is τ , not h , because τ varies over orders of magnitude across network technologies while h varies over single digits.

(f) Risks and misuse. The main misuse path is treating $\Theta = 1/(h\tau)$ as a certified safety bound before the transition is demonstrated. A deployment that sizes its topology to $\lambda < \Theta$ on the strength of this paper alone would be relying on an unvalidated convention (Assumption A1) and unmodeled effects (jitter, loss, coupling). Until the simulation and experimental program confirms the threshold location and its sharpness, the correct use of the result is as a screening heuristic that flags candidate architectures for empirical testing, not as a guarantee.

8. Conclusion

This paper proposed the Network Tracing Capacity $\Theta = 1/(h\tau)$ as a topology-level bound on the instability rates a communication network can propagate for a scalar unstable mode under delayed feedback, derived the bound and its consequences with explicit arithmetic — a five-hop, 10 ms-per-hop topology yields $\Theta = 20 \text{s}^{-1}$, placing a 35s^{-1} mode beyond the horizon at hop 3 of 5 with a per-delay growth factor of $e^{1.75} \approx 5.75$, while reducing the

per-hop delay to 2 ms restores $\Theta = 100 \text{ s}^{-1}$ — and constructed the structural mapping from (instability rate, delay, hop count) onto (expansion rate, signal speed, metric scale). The analogy to a relativistic event horizon is supported at the level of mechanism and operational criteria: a finite propagation rate interacting with exponential growth creates a sharp boundary beyond which information about the growing quantity cannot reach the observer in usable form, a pattern the Event Horizon Telescope literature documents on the physical side, from the stable ring geometry across 2017–2021 [5] to the pipeline-dependent failure of inference near the boundary [2]. What the paper does not establish is that the threshold constant is physical rather than conventional, that the bound survives stochastic delay, packet loss, or vector plants, or that the transition occurs in practice; these are the explicit targets of the falsification test and simulation program laid out in Sections 6 and 7. If the transition is confirmed near $\lambda\tau_d = 1$ and its sharpness survives jitter, the result converts latency budgets into hop-count-aware instability-rate budgets and gives networked control design its first topology-level invariant; if it is refuted, the arithmetic remains as a screening heuristic and the horizon interpretation is retired. Either outcome is decidable with a simulation that fits on a laptop.

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 15 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] Observing---and Imaging---Active Galactic Nuclei with the Event Horizon Telescope. arXiv:1607.03034v1. <https://arxiv.org/abs/1607.03034v1> [2] First M87 Event Horizon Telescope Results. IX. Detection of Near-horizon Circular Polarization. arXiv:2311.10976v1. <https://arxiv.org/abs/2311.10976v1> [3] First Sagittarius A* Event Horizon Telescope Results. III: Imaging of the Galactic Center Supermassive Black Hole. arXiv:2311.09479v1. <https://arxiv.org/abs/2311.09479v1> [4] Mid-Range Science Objectives for the Event Horizon Telescope. arXiv:2410.02986v1. <https://arxiv.org/abs/2410.02986v1> [5] Horizon-scale variability of M87* from 2017--2021 EHT observations. arXiv:2509.24593v1. <https://arxiv.org/abs/2509.24593v1> [6] First Sagittarius A* Event Horizon Telescope Results. II. EHT and Multi-wavelength Observations, Data Processing, and Calibration.

arXiv:2311.08679v1. <https://arxiv.org/abs/2311.08679v1> [7] First M87 Event Horizon Telescope Results. VIII. Magnetic Field Structure near The Event Horizon. arXiv:2105.01173v2. <https://arxiv.org/abs/2105.01173v2> [8] First Sagittarius A* Event Horizon Telescope Results. VI: Testing the Black Hole Metric. arXiv:2311.09484v1. <https://arxiv.org/abs/2311.09484v1> [9] QNFO: Comprehensive Technical Framework for Network Isomorphism [10] QNFO: Boundary-Crossing Paths as a Local, Computable Bound on Network Emergence [11] QNFO: Stability Compiler