Speed, Precision, and Dissipation Bounds for Boundary Maintenance in Finite Distinction Systems: Finite Update Windows, Bottleneck Throughput, and Resource-Ledger Exit Theorems

Yining Wu
Independent Researcher
yining.wu@alumni.upenn.edu

FDS-P6 develops the speed-precision-dissipation paper in the physical bridge sequence of Finite Distinction Systems. P4 studied coarse-grained anti-recurrence: capacity recovery does not reconstruct distinctions erased by a non-injective bottleneck unless inverse information is carried by side records or an enlarged accounting boundary. O3 studied finite-memory record reuse as an operational entropy/resource channel. P7 studied the complementary case in which task identity is carried by invariant quotients or topological side-ledgers. P6 adds the missing dynamic axis: a boundary distinction must not only be representable or recoverable; it must be maintained within a finite update window, at finite precision, under finite resource input. We formalize boundary-maintenance throughput, bottleneck internal rate, precision demand, confidence cost, correction burden, effective causal reach, externalization relief, invariant-compression relief, and a resource-first dissipation ledger. The central result is a conditional exit theorem: if the rate-distortion demand required for boundary maintenance at tolerance ε over update window τ exceeds the sustainable internal update-throughput capacity, the system must enter at least one exit channel–higher resource/dissipation cost, increased error, latency growth, task relaxation, externalization, invariant compression, resource expansion, or boundary failure. The paper does not replace thermodynamic uncertainty relations, quantum speed limits, stochastic thermodynamics, information thermodynamics, rate-distortion theory, or control under communication constraints. It provides an FDS accounting bridge for real-time survival under finite capacity, physical record turnover, causal bandwidth, and boundary-maintenance loss. Deterministic numerical demonstrations illustrate speed-precision load surfaces, resource ledgers, error floors under finite power, effective causal reachability, externalization break-even, invariant-compression relief, regime diagrams, and a Physical-AI worked example normal form.

Reader contract. This paper is a registered FDS physical-bridge paper. It does not derive all physical speed limits from FDS alone. It does not claim that every update costs exactly $k_{\mathrm{B}}T\ln 2$, nor that speed, precision, and dissipation obey a single universal scalar law in every system. Recent quantum many-body clock models show that precision can scale more favorably than simple linear entropy-cost intuitions suggest, reinforcing why P6 is formulated as an exit-channel accounting framework rather than a universal scalar precision-dissipation law [22]. In P6, "dissipation" is used in a broad physical-ledger sense unless a thermal reservoir calibration is explicitly supplied. The formal exit theorems are throughput/resource theorems; thermodynamic dissipation bounds arise only under additional physical calibration. P6 provides a quantitative audit standard for real-time survival in embodied finite systems. A distinction that cannot be updated, verified, corrected, or acted upon within the relevant update window is not operationally maintained.

Claim-status summary

Table I summarizes the main claims and demotion conditions.

INTRODUCTION

The missing dynamic axis

P4 established a recovery asymmetry: when a non-injective truncation removes task-relevant preimage information from the effective record, later capacity recovery does not reconstruct the erased distinction without side information [2]. P7 established the protected complement: some task identity can be carried by an invariant quotient $Q_{\mathrm{inv}} = q(X)$ rather than by ordinary local memory [4]. O3 established that finite-memory record turnover enters a coupled entropy/resource ledger when physical records must be reused, refreshed, repaired, synchronized, externalized, or erased [3].

P6 adds the missing dynamic axis. A finite system may have enough capacity to represent a distinction and may even know which invariant quotient matters, yet fail if it cannot update, verify, correct, or act on that distinction fast enough and accurately enough to maintain its boundary. A boundary-maintaining system is not a static database. It is a finite physical process operating over an update window τ. The relation to existing theories (rate-distortion, thermodynamic uncertainty relations, quantum speed limits, control theory) is summarized in Table IV; P6 does not replace these theories but

TABLE I. Central FDS-P6 claims, status, and failure or demotion conditions.

Claim	Status	What would weaken or falsify it
Boundary maintenance requires finite update throughput.	Formal FDS claim	A time-varying boundary is maintained without updating, verifying, storing, externalizing, protecting, or acting on any task-relevant distinction.
Speed and precision jointly increase maintenance burden.	Operational bridge	Faster and more precise maintenance is sustained indefinitely at fixed representation and fixed resource input, with no extra dissipation, error, latency, externalization, invariant compression, or failure.
The sustainable internal rate is bottlenecked by sensing, updating, verification, correction, action, and resources.	Bottleneck definition	A system exceeds its slowest internal channel indefinitely without queueing, latency, loss, externalization, or resource expansion.
Correction and verification belong in the resource ledger.	O3-compatible physical bridge	Physical correction, refresh, verification, synchronization, overwrite, and recovery are cost-free under the stated implementation assumptions.
Effective causal update bandwidth limits real-time maintenance.	Physical/engineering bridge	A finite observer integrates arbitrarily distant boundary-relevant information within a finite update window with no latency, no prediction burden, and no effective signal-speed limit.
Externalization and invariant compression are relief channels, not free exits.	P4/P7-compatible bridge	External ledgers or invariant quotients reduce internal demand with no write, synchronization, verification, protection, latency, or boundary-accounting cost.

places them inside a boundary-maintenance accounting problem.

The physical bridge ladder

P6 is placed in the FDS physical bridge ladder as shown in Table II. Its role is not to replace P4, O3, or P7, but to bind them into a throughput constraint.

TABLE II. P6 in the local physical bridge ladder.

Paper	Core axis	P6 relation
P4	recoverability	lost preimages cannot be speed-recovered
O3	record turnover	update channels enter a ledger
P6	throughput	speed, precision, and cost constrain maintenance
P7	protection	quotients reduce maintained detail but not all cost

Speed, precision, and resource ledgers

Speed means how often the boundary-relevant state must be refreshed or acted upon. Precision means how tightly the maintained representation must approximate the task-relevant boundary statistic. Dissipation means the physical resource or entropy ledger associated with sensing, updating, verifying, correcting, synchronizing, externalizing, protecting, and recovering that representation. The central P6 thesis is: finite boundary maintenance requires finite throughput, finite precision, and a finite ledger. At fixed resources, pushing speed and precision together forces an exit channel.

Why effective causal update bandwidth matters

A physically realized system cannot instantaneously integrate all boundary-relevant information. A finite maximum signal velocity implies a finite causal distinguishability-update bound. P6 uses the weaker operational statement: given an effective maximum propagation speed $v_{\max}$, real-time maintenance has a causal reach $R_H(\tau) = v_{\max}\tau$. In engineered or biological systems, $v_{\max}$ may be sensor, actuator, communication, diffusion, or transport limited rather than equal to the relativistic upper bound c.

FDS BACKGROUND

Active finite distinction systems

The FDS core object is

$$S = (X, E, B, M, Y, A, U, \pi, \ell, \Phi, \mathcal{P}, \tau), \tag{1}$$

where X is internal state, E environment, B boundary, M memory/model state, Y observation channel, A action space, U update rule, π finite projection, ℓ boundary-maintenance loss, Φ finite resource budget, $\mathcal{P}$ perturbation/pruning family, and τ update timescale [1]. P6 focuses on U, π, ℓ, Φ, Y, M, and τ.

Capacity demand and update windows

Let Ψ be a task family and let $R^{(\tau)}_{\min}(\varepsilon;\Psi)$ be the minimum number of bits over update window τ needed to encode a task-relevant statistic to tolerance ε. The capacity deficit is

$$\Delta_\varepsilon(\tau) = R^{(\tau)}_{\min}(\varepsilon;\Psi) - C_{\text{eff}}. \tag{2}$$

P6 turns this into a rate:

$$\mathcal{B}_{\varepsilon,\tau} := \frac{R^{(\tau)}_{\min}(\varepsilon;\Psi)}{\tau}. \tag{3}$$

A distinction that is representable in total storage may still be unmaintainable if $\mathcal{B}_{\varepsilon,\tau}$ exceeds the rate at which the system can sense, update, verify, correct, and act.

DEFINITIONS

Definition 1 (Boundary-maintenance task). *Let $Z_t = \psi(E_t, B_t)$ be a pre-registered task-relevant boundary statistic. A system maintains the task at tolerance ε over update window τ if*

$$\mathbb{E}[\ell(Z_t, \widehat{Z}_t)] \le \varepsilon \tag{4}$$

within the window and under the stated perturbation family.

Definition 2 (Maintenance speed). *For clocked systems, the maintenance speed is*

$$\nu_{\text{upd}} = \frac{1}{\tau}. \tag{5}$$

For event-driven systems, one may use

$$\nu_{\text{upd}}(t) = \frac{N_{\text{task}}(t, t+\tau)}{\tau}, \tag{6}$$

where N_{task} counts task-relevant update events.

Definition 3 (Precision and confidence). *Precision is encoded by tolerance ε and confidence $1-\delta$:*

$$\mathbb{P}\{D(\widehat{Z}_t, Z_t) \le \varepsilon\} \ge 1-\delta. \tag{7}$$

Smaller ε and smaller δ increase the required distinction throughput and verification burden.

Definition 4 (Boundary-maintenance throughput). *The throughput burden of maintaining task family Ψ at tolerance ε over update window τ is Eq. (3). It has units of task-relevant distinction-bits per unit time.*

Definition 5 (Correction burden). *Let $p_{\text{err}}(\tau,\varepsilon,\delta)$ be a boundary-relevant error probability and C_{repair} the expected repair or correction cost after such an error. The expected correction burden is*

$$\bar{\mathcal{C}}_{\text{corr}} := \mathbb{E}[p_{\text{err}}\, C_{\text{repair}}]. \tag{8}$$

More detailed implementations may replace this by a stochastic thermodynamic, control, or recovery-cost model.

Definition 6 (Internal sustainable rate). *The internally sustainable maintenance rate is the bottleneck*

$$C^{\text{rate}}_{\text{int}} = \min\{C^{\text{rate}}_{\text{sense}}, C^{\text{rate}}_{\text{update}}, C^{\text{rate}}_{\text{verify}}, C^{\text{rate}}_{\text{correct}}, C^{\text{rate}}_{\text{act}}, C^{\text{rate}}_{\text{resource}}\}. \tag{9}$$

The resource-limited term may be represented in a model class as

$$C^{\text{rate}}_{\text{resource}} = \frac{\dot{\Phi}_{\text{in}}}{\kappa_{\varepsilon,\delta,\tau}}, \tag{10}$$

where $\kappa_{\varepsilon,\delta,\tau}$ is the implementation-dependent resource cost per maintained distinction-rate at the specified tolerance, confidence, and update window. This is not a universal constant. When the system faces a transient high-throughput demand, increasing $\dot{\Phi}_{\text{in}}$ alone may not prevent boundary collapse because the physical limit on the internal update rate $C^{\text{rate}}_{\text{int}}$ creates a nonlinear saturation: beyond a certain point, additional energy cannot be converted into higher update throughput due to bottleneck constraints on sensing, actuation, or verification speed.

Definition 7 (P6 resource ledger). *The P6 resource ledger is an audit partition*

$$\begin{aligned}\dot{\mathcal{R}}_{\text{P6}} =& \dot{\mathcal{R}}_{\text{sense}} + \dot{\mathcal{R}}_{\text{update}} + \dot{\mathcal{R}}_{\text{verify}} + \dot{\mathcal{R}}_{\text{correct}} \\ &+ \dot{\mathcal{R}}_{\text{sync}} + \dot{\mathcal{R}}_{\text{externalize}} + \dot{\mathcal{R}}_{\text{protect}}.\end{aligned} \tag{11}$$

When thermodynamic calibration is available, a resource ledger may be converted into an entropy ledger, for example $\dot{\mathcal{R}}_{\text{P6}} \to T_{\text{eff}}\dot{\Sigma}_{\text{P6}}$ under appropriate reservoir assumptions. Without such calibration it remains a resource, power, latency, or cost ledger.

MAIN BOUNDS

Theorem 1 (Throughput deficit exit theorem). *Consider a physically realized active FDS maintaining a boundary task Ψ at tolerance ε over update window τ, with finite effective capacity C_{eff}, finite resource input $\dot{\Phi}_{\text{in}}$, and a fixed accounting boundary $\mathcal{A}$. If*

$$\frac{R^{(\tau)}_{\min}(\varepsilon;\Psi)}{\tau} > C^{\text{rate}}_{\text{int}}, \tag{12}$$

then sustained maintenance requires at least one exit channel: increased resource/dissipation ledger, increased error, increased latency, task relaxation, externalization, invariant compression, resource expansion, or boundary-maintenance failure. Recent work on maximum dissipation speed and fluctuation speed limits provides nearby model-class examples in which time, dissipation, and observable precision are explicitly constrained [19, 20].

Proof. The task requires $R^{(\tau)}_{\min}(\varepsilon;\Psi)$ task-relevant bits over each update window, hence a throughput burden $\mathcal{B}_{\varepsilon,\tau}$. By definition, $C^{\mathrm{rate}}_{\mathrm{int}}$ is the slowest internal channel among sensing, update, verification, correction, action, and resource-limited throughput. If $\mathcal{B}_{\varepsilon,\tau}$ exceeds this bottleneck, the missing throughput must be supplied by some other channel or the demand must be reduced. The possibilities are: pay additional resource cost, accept more error, increase latency by lengthening the effective update window, reduce the task demand, move records to an external ledger, replace raw distinctions by a lower-rate invariant quotient, expand resources, or fail the boundary task. If none occurs, the system lacks the throughput needed by the definition of $R^{(\tau)}_{\min}(\varepsilon;\Psi)$, contradicting sustained maintenance at tolerance ε. □

Full P6 ledger demand

The full P6 demand includes rate-distortion throughput, confidence burden, correction burden, and causal prediction burden. Here $C^{\mathrm{rate}}_{\mathrm{sust}}$ denotes the sustainable maintained-throughput rate under the fixed accounting boundary, including the internal bottleneck and any baseline resource-supported maintenance capacity before relief channels are applied. In the minimal internal case, $C^{\mathrm{rate}}_{\mathrm{sust}} = C^{\mathrm{rate}}_{\mathrm{int}}$:

$$D_{\mathrm{P6}} = \frac{R^{(\tau)}_{\min}(\varepsilon;\Psi)}{\tau} + \lambda_\delta \frac{\log_2(1/\delta)}{\tau} + \lambda_c \frac{\bar{\mathcal{C}}_{\mathrm{corr}}}{\tau} + D_{\mathrm{causal}}. \tag{13}$$

The logarithmic confidence term is not asserted as a universal equality. It is a normal-form reliability cost motivated by coding, hypothesis testing, and repeated verification settings in which reducing failure probability from $O(1)$ to δ requires information or verification effort scaling at least logarithmically in $1/\delta$. Here D_{causal} denotes the rate-equivalent prediction, delay, or error burden induced when task-relevant boundary information lies outside the effective causal reach $R_H(\tau) = v_{\max}\tau$.

Theorem 2 (Full P6 ledger exit theorem)**.** *Let $R^{\mathrm{relief}}_{\mathrm{inv}}$ and $R^{\mathrm{relief}}_{\mathrm{ext}}$ denote relief supplied by invariant compression and externalization after subtracting their coupled protection, synchronization, verification, retrieval, and latency ledgers. If*

$$D_{\mathrm{P6}} > C^{\mathrm{rate}}_{\mathrm{sust}} + R^{\mathrm{relief}}_{\mathrm{inv}} + R^{\mathrm{relief}}_{\mathrm{ext}}, \tag{14}$$

then the same boundary-maintenance regime cannot persist without higher resource/dissipation ledger, higher error, higher latency, task relaxation, accounting-boundary expansion, invariant compression, resource expansion, or failure.

Proof. Equation (13) is the task-relevant demand rate for maintenance under the chosen tolerance, confidence, correction model, and causal reach. The right-hand side is the sustainable rate plus relief channels that actually reduce the maintained burden. If the demand exceeds this total, then either the ledger must increase, the representation/task must change, latency or error must grow, the accounting boundary must expand, or the boundary condition is not maintained. Otherwise the system would be maintaining a demand rate greater than the available rate without a source of the missing throughput. □

Normal-form resource inequality

A useful model-class bridge is

$$\dot{\mathcal{R}}_{\mathrm{P6}} \geq \alpha \frac{R^{(\tau)}_{\min}(\varepsilon;\Psi)}{\tau} + \beta \frac{1}{\tau}\log_2\frac{1}{\delta} + \gamma \frac{\bar{\mathcal{C}}_{\mathrm{corr}}}{\tau} - \mathcal{R}^{\mathrm{relief}}_{\mathrm{inv}} - \mathcal{R}^{\mathrm{local\ relief}}_{\mathrm{ext}} + \mathcal{R}^{\mathrm{coupled}}_{\mathrm{ext}} + \mathcal{R}_{\mathrm{protect}}. \tag{15}$$

Here α is a cost per maintained distinction-rate, β a confidence or verification cost, γ a correction coefficient, $\mathcal{R}^{\mathrm{relief}}_{\mathrm{inv}}$ is rate reduction from invariant compression, $\mathcal{R}^{\mathrm{local\ relief}}_{\mathrm{ext}}$ is local relief from externalization, $\mathcal{R}^{\mathrm{coupled}}_{\mathrm{ext}}$ is the coupled external ledger, and $\mathcal{R}_{\mathrm{protect}}$ is the cost of maintaining a protected quotient. Equation (15) is a normal form, not a universal identity.

Corollary 1 (No zero-latency, zero-error, zero-cost limit)**.** *At fixed finite capacity and fixed finite resource input, a finite active boundary cannot simultaneously take*

$$\tau \to 0, \qquad \varepsilon \to 0, \qquad \delta \to 0 \tag{16}$$

while maintaining the same task without an exit channel.

Corollary 2 (Latency-precision tradeoff under finite resources)**.** *For a model class with sustainable rate $C^{\mathrm{rate}}_{\mathrm{sust}}$, if $R_{\min}(\varepsilon;\Psi) \sim d\log_2(1/\varepsilon)$ and confidence/correction terms are fixed, then the update window must satisfy approximately*

$$\tau \gtrsim \frac{d\log_2(1/\varepsilon)}{C^{\mathrm{rate}}_{\mathrm{sust}}}. \tag{17}$$

Thus, at fixed sustainable rate, tighter precision requires longer latency unless relief or extra resources are supplied.

PRECISION-SPEED TRADEOFF

For many regular task classes, the rate-distortion demand scales as

$$R_{\min}(\varepsilon; \Psi) \sim d \log_2 \frac{1}{\varepsilon}, \tag{18}$$

where d is an effective task dimension. Then

$$\mathcal{B}_{\varepsilon,\tau} \sim \frac{d}{\tau} \log_2 \frac{1}{\varepsilon}. \tag{19}$$

This captures the most basic P6 load: faster updates and tighter precision multiply. Confidence adds a verification term proportional to $\tau^{-1} \log_2(1/\delta)$ in many normal-form reliability settings.

Proposition 1 (Finite-resource precision floor)**.** *Suppose a model class has a maintenance ledger lower bound*

$$\dot{\mathcal{R}} \geq a \frac{d}{\tau} \log_2 \frac{1}{\varepsilon} + b \frac{1}{\tau} \log_2 \frac{1}{\delta}, \tag{20}$$

with $a, b > 0$. If $\dot{\mathcal{R}} \leq \dot{\mathcal{R}}_{\max}$ and δ is fixed, then the achievable precision obeys

$$\varepsilon \geq 2^{-\frac{\tau}{ad}\left(\dot{\mathcal{R}}_{\max} - b\tau^{-1} \log_2(1/\delta)\right)} \tag{21}$$

whenever the exponent is positive. Otherwise the requested speed-confidence combination is infeasible in that model class.

EFFECTIVE CAUSAL UPDATE BANDWIDTH

Finite causal reach

Given an effective maximum propagation or integration speed $v_{\max}$, boundary-relevant information from distance R cannot be integrated into the present update if $R > v_{\max}\tau$. Define the effective causal reach in one update window:

$$R_H(\tau) = v_{\max}\tau. \tag{22}$$

For relativistic causal bounds, $v_{\max} \leq c$ [17]. In engineered or biological systems, $v_{\max}$ may be sensor, actuator, communication, diffusion, transport, or neural-conduction limited. P6 requires only finite effective causal reach. Recent feedback-emulation results sharpen the data-rate perspective by asking not only whether a system can be stabilized, but what rate is needed to emulate a desired closed-loop behavior [23].

Maintenance reachability

Let $\Psi(A)$ be the set of task distinctions supported on a boundary region A. During update window τ, the maintainable task family is restricted to the causal reach:

$$\Psi_{\text{maintainable}}(\tau) \subseteq \Psi(A_{\text{reach}}(\tau)). \tag{23}$$

If a task requires boundary-relevant information outside $A_{\text{reach}}(\tau)$, the system must predict, externalize, delay, relax the task, or fail. This is not a claim about faster-than-light impossibility in general; it is a boundary-maintenance accounting statement.

LEDGER CHANNELS

Update and verification

An update channel carrying B_{upd} bits per window has a resource-rate proxy

$$\dot{\mathcal{R}}_{\text{update}} \sim \eta_{\text{upd}} \frac{B_{\text{upd}}}{\tau}. \tag{24}$$

A verification channel enforcing failure probability δ carries

$$\dot{\mathcal{R}}_{\text{verify}} \sim \eta_{\text{verify}} \frac{\log_2(1/\delta)}{\tau}. \tag{25}$$

These proxies are not universal thermodynamic laws. They are audit terms to be replaced by implementation-specific measurement when available.

Correction

If errors occur with probability p_{err} per window and the mean correction cost is C_{repair}, then

$$\dot{\mathcal{R}}_{\text{correct}} \sim \frac{p_{\text{err}} C_{\text{repair}}}{\tau}. \tag{26}$$

Thermodynamic uncertainty relations and speed-precision tradeoffs in stochastic thermodynamics provide mature model classes in which precision, current fluctuations, speed, and dissipation are formally related [12–14]. Kinetic uncertainty relations further show that response precision can be bounded by dynamical activity in Markovian nonequilibrium systems, supporting the P6 treatment of verification and correction as ledger channels [21].

Externalization

Externalization can reduce the internal fast-update load, but it creates a coupled ledger:

$$\begin{aligned}\dot{\mathcal{R}}_{\text{ext}} =& \dot{\mathcal{R}}_{\text{write}} + \dot{\mathcal{R}}_{\text{sync}} + \dot{\mathcal{R}}_{\text{retrieve}} \\ &+ \dot{\mathcal{R}}_{\text{verify}} + \dot{\mathcal{R}}_{\text{protect}} + \dot{\mathcal{R}}_{\text{latency}}.\end{aligned} \tag{27}$$

This is the dynamic analogue of the P4 side-record audit: exact recovery or high-speed maintenance can be shifted to an external log, but not made free [2].

Invariant compression

If task identity factors through an invariant quotient,

$$Q_{\text{inv}} = q(X), \qquad V = h(Q_{\text{inv}}), \tag{28}$$

then the maintained rate can drop from a raw representation rate to a quotient rate:

$$\Delta R_{\text{inv}} = R_{\text{raw}}(\varepsilon) - R_q(\varepsilon). \tag{29}$$

P7 emphasizes that the protection ledger must still be counted: topological or invariant persistence relocates cost to boundary, drive, verification, or protection maintenance rather than deleting it [4].

WORKED EXAMPLE: HIGH-SPEED EMBODIED MAINTENANCE

Consider a mobile embodied agent maintaining a collision-avoidance boundary. The task statistic Z_t encodes nearby obstacle position and relative velocity, the update window τ is the control period, and ε is the maximum admissible tracking or clearance error. A remote high-latency model controller (including cloud LLM-style controllers) may have high semantic precision but high latency, communication cost, and verification overhead. A local reflex or morphology-based controller may have lower semantic resolution but much smaller τ, lower ledger cost, and direct actuation. P6 predicts that, under tight physical update windows, the second system can dominate survival even if its representation is less detailed.

In P6 language, the cloud controller pays a large D_{causal} and verification/externalization ledger, while the local controller reduces τ and shifts task identity into local feedback, morphology, or invariant carriers. For physical survival tasks, lower representational precision can dominate when it reduces latency and ledger cost. This is not an argument against large models. It is a statement that boundary maintenance is constrained by throughput, not capacity alone.

NUMERICAL DEMONSTRATIONS

The simulations are deterministic normal-form demonstrations generated by `code/generate_results.py`. They illustrate definitions and exit channels rather than fitting empirical data. The regime diagram is a synthetic classification map generated from threshold rules specified in the reproducibility code. It illustrates exit-channel ordering and should not be read as an empirical phase diagram.

TABLE III. Core demonstration parameters. Full parameter values are exported to `data/simulation_parameters.csv`.

Parameter	Value / role
d	8, task-dimension proxy
δ	10^{-3}, confidence target
Resource thresholds	$120, 220, 360$ in normalized units
$v_{\max}$	$1.0, 0.25, 0.06$ normalized reach speeds
Energy calibration	optional Landauer floor at 300 K

EXPERIMENTAL AND ENGINEERING PROTOCOLS

Protocol 1 (Finite update-window stress test)**.** *Hold task demand and precision target fixed while reducing τ. Measure error rate, latency, power, reset/overwrite rate, correction count, externalization count, and boundary loss. P6 predicts ledger increase, error increase, latency, externalization, invariant compression, task relaxation, or failure.*

Protocol 2 (Precision tightening test)**.** *Hold update window fixed while lowering ε and δ. Measure verification, correction, resource cost, and latency. P6 predicts increased throughput burden and a finite-resource error floor.*

Protocol 3 (Invariant-compression test)**.** *Compare raw high-dimensional representation, externalized record, and quotient/invariant representation. Measure task loss, ledger cost, update speed, and robustness under perturbation. A successful quotient reduces rate burden without exceeding the protection ledger.*

Protocol 4 (Physical AI maintenance benchmark)**.** *For an embodied agent, stress moving-target tracking, balance control, navigation, manipulation under sensor damage, or long-horizon memory-constrained action. Report $\mathcal{B}_{\varepsilon,\tau}$, ledger cost, boundary loss, error rate, recovery time, and externalization/protection channels.*

FALSIFICATION AND DEMOTION CONDITIONS

P6 is weakened if a finite physically realized system maintains arbitrarily fast and arbitrarily precise boundary tasks at fixed capacity and fixed resource input with no increase in resource ledger, error, latency, externalization, invariant compression, task relaxation, or hidden accounting-boundary expansion. A model-class inequality such as Eq. (15) can fail without falsifying the FDS exit theorem; it should then be demoted to a narrow toy model. If hidden external records or uncounted resources are discovered, the accounting boundary must be expanded. If a protected quotient carries the task iden-

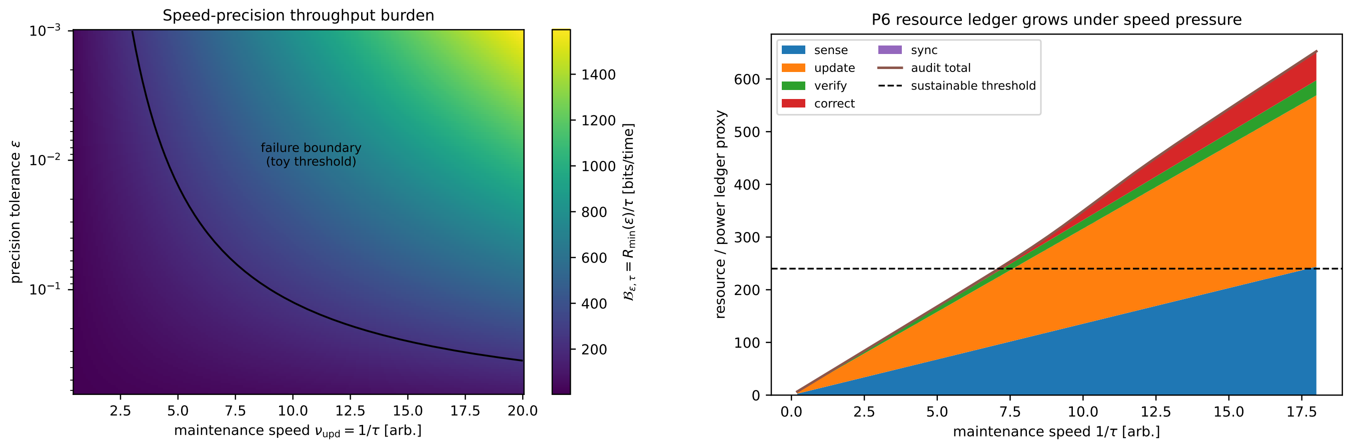

FIG. 1. Left: speed-precision throughput burden $\mathcal{B}_{\varepsilon,\tau} = R_{\min}(\varepsilon)/\tau$ for a task dimension proxy $d = 8$. The contour is a toy failure boundary. Right: P6 resource ledger decomposition under increasing speed pressure. Channels are illustrative audit classes, not universal thermodynamic mechanisms.

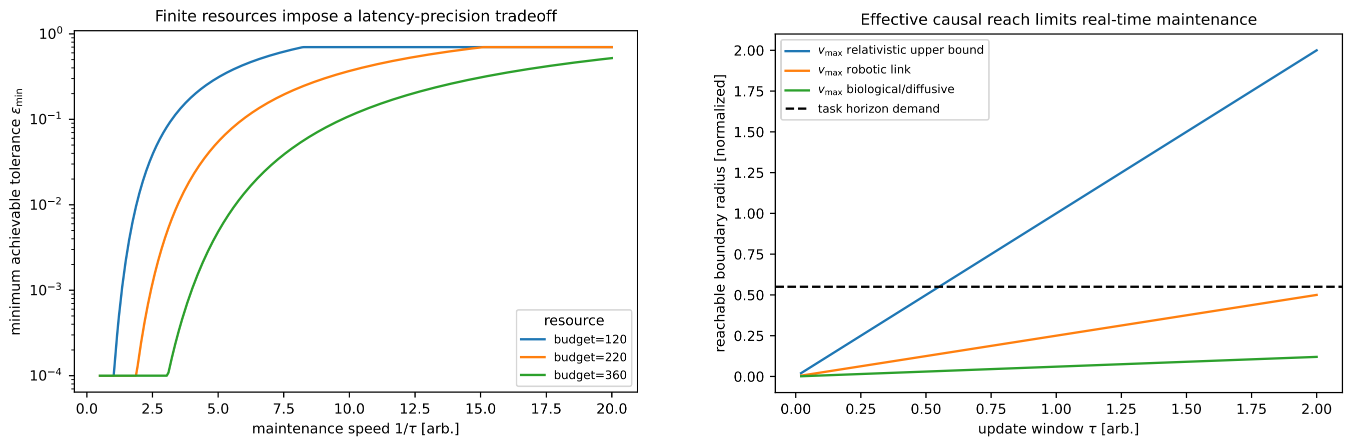

FIG. 2. Left: finite resources impose a speed-dependent precision floor. At high update speed, the same resource budget cannot support arbitrarily small error tolerance. Right: effective causal reach $R_H = v_{\max}\tau$ shrinks with the update window. Boundary-relevant information outside the reach must be predicted, externalized, delayed, or abandoned. The protection-breaking load occurs where the causal horizon can no longer cover the boundary region requiring maintenance.

tity, the case belongs to the P7 relief channel rather than being a counterexample to P6.

RELATION TO EXISTING THEORY

P6 uses mature theories as implementation bridges rather than replacing them: P6 does not replace any of these theories. It packages their common operational implication inside FDS: boundary maintenance is throughput-constrained.

GLOSSARY

Term	P6 meaning
Speed	update frequency $1/\tau$
Precision	distortion tolerance ε
Confidence	success probability $1-\delta$
Throughput	task-relevant distinction rate $\mathcal{B}_{\varepsilon,\tau}$
Ledger	audited resource or entropy-cost channels
Exit channel	cost, error, latency, externalization, quotient, relaxation, or failure

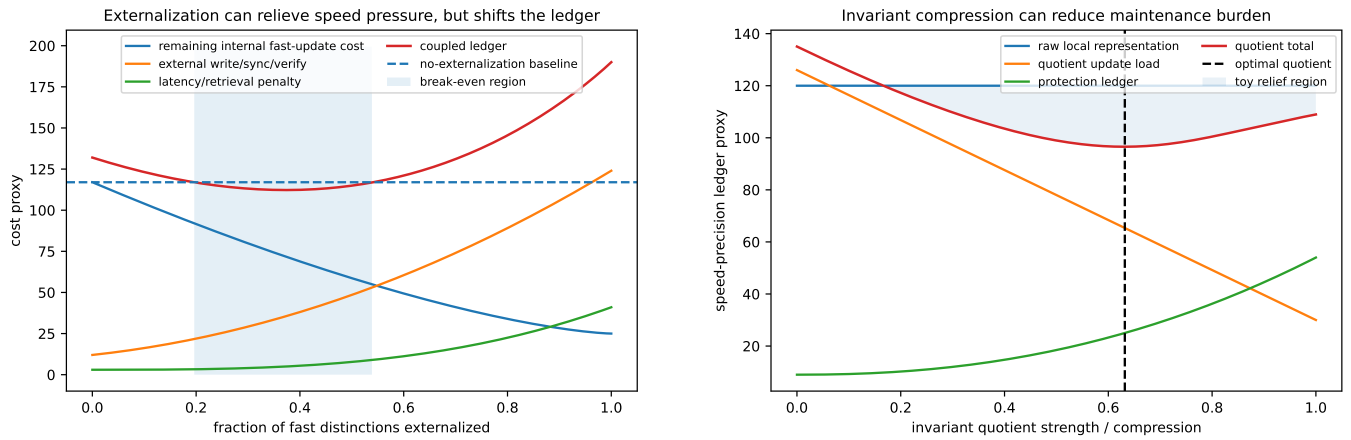

FIG. 3. Left: externalization can relieve internal speed pressure but shifts cost into writing, synchronization, verification, retrieval, latency, and protection. Right: invariant compression can reduce update burden when task identity factors through a quotient, but the protection ledger and task-loss penalty create an optimum rather than monotonic improvement. The exported CSV includes a toy Landauer-calibrated power floor at 300 K; it is not used as a universal heat claim.

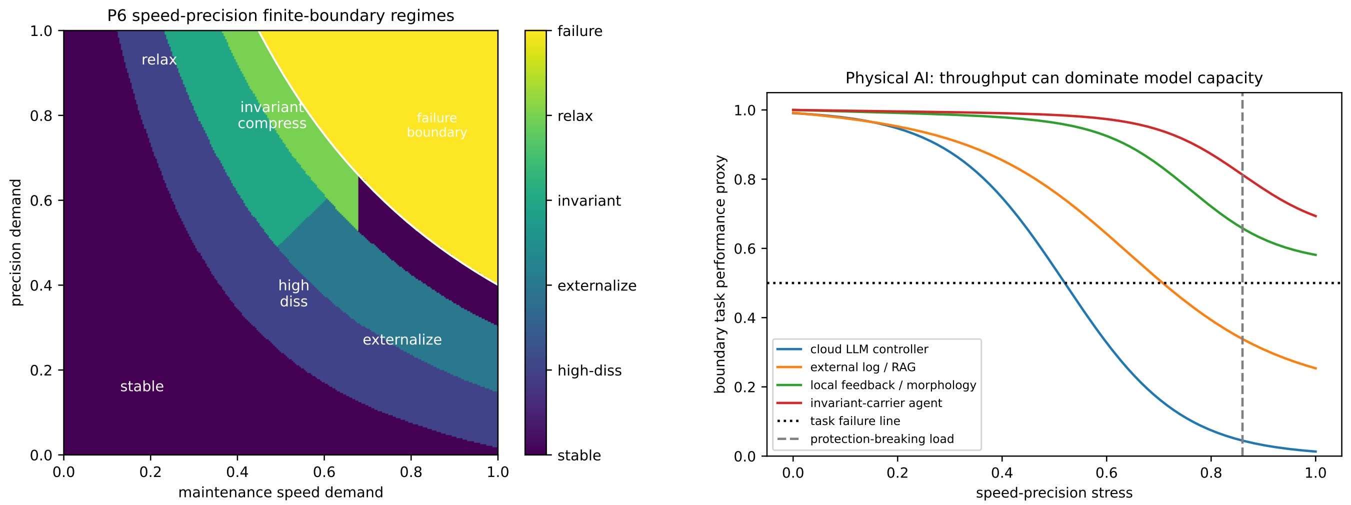

FIG. 4. Left: P6 speed-precision finite-boundary regimes. Increasing speed and precision moves a finite system from stable maintenance into high dissipation, externalization, invariant compression, relaxation, or failure. Region boundaries are synthetic threshold rules. Right: Physical-AI worked example normal form. Local feedback and invariant carriers can outperform a remote high-latency model controller under high speed-precision stress because they reduce latency and ledger cost.

DESIGN IMPLICATION: BOUNDARY MAINTENANCE UNDER SPEED-PRECISION PRESSURE

This section is not part of the formal proof. It extracts the FDS design rule implied by P6. A physical AI agent is not limited only by model size. It is limited by how fast, how precisely, and how cheaply it can maintain task-relevant boundaries. The design rule is: do not maximize memory alone; optimize the speed-precision-dissipation ledger. High-speed tasks should use invariants, reflexive control loops, morphology, and local closed-loop policies. Slow exact recovery can use external logs. When causal update is delayed, predictive control is required. When failure cost is high, verification should increase. When speed matters more than detail, precision should relax. P6 thereby provides a benchmark language for Physical AI: report maintenance throughput, not only accuracy.

CONCLUSION

P4 showed that lost distinctions do not return merely because capacity returns. P7 showed that some task identity can survive local forgetting through invariant quotients. O3 showed that finite record reuse enters an entropy/resource ledger. P6 completes this portion of the physical spine by adding speed. A finite system must

TABLE IV. P6 relation to existing theories.

Existing theory	What it bounds	P6 use
Rate-distortion theory	bits needed for precision	task demand $R_{\min}(\varepsilon)$ [5–7]
Information thermodynamics	erasure, measurement, feedback, memory	physical ledger terms [8–11]
Thermodynamic uncertainty relations	precision versus dissipation	stochastic implementation bridge [12–14]
Quantum speed limits	minimum transformation time	physical update-time constraints [15, 16]
Data-rate/control theory	feedback control under channel limits	communication-limited boundary maintenance [18,
Finite-time dissipation speed limits	time vs entropy/resource rate	P6 speed-cost bridge [19, 20]
Response / kinetic uncertainty relations	activity vs response precision	verification/correction burden [21]
FDS-P6	finite boundary update throughput	exit-channel accounting under finite capacity

maintain its boundary not only with enough representation, but within a finite update window, at finite precision, and under finite resource budgets.

P6 identifies the dynamic cost of being finite. Pushing speed and precision together forces dissipation, error, latency, externalization, invariant compression, task relaxation, or failure. Boundary maintenance is therefore not a static property of representation. It is a throughput-constrained physical process. Throughout P6, "dissipation" is used as generalized resource turnover cost; thermodynamic dissipation is a special case that arises when thermal reservoir calibration is available. For Physical AI, this means that durable agency is not simply larger memory or larger models. It is the ability to preserve task-relevant boundaries under the speed, precision, and dissipation constraints of the world.

[1] Y. Wu, "Active Finite Distinction Systems: A Formal Core for Boundary Maintenance under Finite Capacity," Zenodo (2026), doi:10.5281/zenodo.20158923.

[2] Y. Wu, "Coarse-Grained Anti-Recurrence and Informational Hysteresis in Finite Memory Systems: Lost Preimages, Side Records, and Capacity-Recovery Asymmetry," Zenodo (2026), doi:10.5281/zenodo.20265065.

[3] Y. Wu, "Boundary Maintenance and the Second Law under Finite Memory: Irreversible Record Reuse, Entropy Ledgers, and Operational Time Arrows," Zenodo (2026), doi:10.5281/zenodo.20255129.

[4] Y. Wu, "Topological Obstruction to Forgetting in Finite Distinction Systems: Quotient Invariants, Non-Hermitian Skin Effects, and Topological Side-Ledgers," Zenodo (2026), doi:10.5281/zenodo.20265386.

[5] C. E. Shannon, "A mathematical theory of communication," Bell Syst. Tech. J. **27**, 379–423 and 623–656 (1948).

[6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. (Wiley, Hoboken, 2006).

[7] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, 2003).

[8] R. Landauer, "Irreversibility and heat generation in the computing process," IBM J. Res. Dev. **5**, 183–191 (1961), doi:10.1147/rd.53.0183.

[9] C. H. Bennett, "The thermodynamics of computation - a review," Int. J. Theor. Phys. **21**, 905–940 (1982).

[10] U. Seifert, "Stochastic thermodynamics, fluctuation theorems and molecular machines," Rep. Prog. Phys. **75**, 126001 (2012).

[11] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, "Thermodynamics of information," Rev. Mod. Phys. **87**, 45–77 (2015), doi:10.1103/RevModPhys.87.45.

[12] A. C. Barato and U. Seifert, "Thermodynamic uncertainty relation for biomolecular processes," Phys. Rev. Lett. **114**, 158101 (2015), doi:10.1103/PhysRevLett.114.158101.

[13] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, "Dissipation bounds all steady-state current fluctuations," Phys. Rev. Lett. **116**, 120601 (2016), doi:10.1103/PhysRevLett.116.120601.

[14] J. M. Horowitz and T. R. Gingrich, "Thermodynamic uncertainty relations constrain non-equilibrium fluctuations," Nat. Phys. **16**, 15–20 (2020), doi:10.1038/s41567-019-0702-6.

[15] L. Mandelstam and I. Tamm, "The uncertainty relation between energy and time in non-relativistic quantum mechanics," J. Phys. (USSR) **9**, 249–254 (1945).

[16] N. Margolus and L. B. Levitin, "The maximum speed of dynamical evolution," Physica D **120**, 188–195 (1998), doi:10.1016/S0167-2789(98)00054-2.

[17] A. Einstein, "Zur Elektrodynamik bewegter Körper," Ann. Phys. **322**, 891–921 (1905).

[18] S. Tatikonda and S. Mitter, "Control under communication constraints," IEEE Trans. Autom. Control **49**, 1056–1068 (2004), doi:10.1109/TAC.2004.831187.

[19] S. Das and J. R. Green, "Maximum speed of dissipation," Phys. Rev. E **109**, L052104 (2024), doi:10.1103/PhysRevE.109.L052104.

[20] R. Hamazaki, "Speed limits to fluctuation dynamics," Commun. Phys. **7**, 361 (2024), doi:10.1038/s42005-024-01851-y.

[21] K. Liu and J. Gu, "Dynamical activity universally bounds precision of response in Markovian nonequilibrium systems," Commun. Phys. **8**, 62 (2025), doi:10.1038/s42005-025-01982-w.

[22] F. Meier, Y. Minoguchi, S. Sundelin, T. J. G. Apollaro, P. Erker, S. Gasparinetti, and M. Huber, "Precision is not limited by the second law of thermodynamics," Nat. Phys. **21**, 1147–1152 (2025), doi:10.1038/s41567-025-02929-2.

[23] G. S. Vicinansa and G. Nair, "Minimum Data-Rate for Emulating a Linear Feedback System," IFAC-PapersOnLine **58**(17), 350–355 (2024),

doi:10.1016/j.ifacol.2024.10.194.