

Sampled Tracking with a Changing Roster

Bounds for a nominal reference controller

Guillaume Bolduc
HORTRAME Systems

7 October 2026 | Internal preprint, source version 0.6

Abstract

We study a sampled reference controller on a finite set of nominal double integrators with changing logical participation. A correct reference service, coherent delayed physical tables, a common clock, and a constrained translation latch are specified as resources. An error estimate accounts for disturbance integrals, goal changes, admission state, snapshot age, and projection error. A weighted event budget gives a tracking bound without requiring participation eventually to become fixed. Separate convex-cell assumptions give continuous separation, speed, and acceleration bounds for all bodies, including inactive bodies. Under a supplied finite service discipline, a delayed-observation gate permits repeated logical reinitialization and readmission of the same bodies without accumulating bodies or workspace. Uniform operational storage additionally requires bounded message production, frame sizes, lifetime, and serviced queues. A convergent quiet tail gives eventual latch freeze under the stated contraction condition. These are conditional results for fixed task coordinates, reserved operating cells, exact local feedback, and correct services; numerical and physical realization remain separate questions.

Status. This source revision responds to an additional source-bound review of v0.4 following the separately frozen ENTAIL and FOLIUM reviews. Their shared sources and tooling do not constitute independent human validation. The matched v0.5 source/PDF pair remains the most recent compiled edition; PDF work on this source revision is owner-paused. Human acceptance, numerical certification, and physical realization remain open.

1 Problem, source frame, and position relative to prior work

A network that changes forever generally has no final consensus target. A physically persistent system also cannot consume a new body or a new parking berth at every change. The question studied here is therefore an infinite-execution question: which information, control, and physical conditions suffice to keep a finite reusable team safe and quantitatively close to its current task while logical membership continues to change?

The supplied foundation [1] distinguishes consensus from restoration, identifies history-dependent endpoints, and constructs safe finite replacement under separately stated contracts. Its convergence claims require a final fixed-support epoch, and its physical parking construction uses each body once. The disturbance paper [2] separates common translation from relative geometry. The latch addendum [3] proves nominal fixed-cohort convergence through an energy contraction and a position projection. The admission addendum [4] removes a newcomer’s physical residual from incumbent feedback and identifies a stale-history release condition. The finite-history paper [5] counts retained information needed for a specified later query. None of these statements by itself proves their simultaneous operation under endless membership turnover.

This article makes four deliberate model changes. It uses a fixed task alphabet, rather than creating a new task coordinate for every incarnation. It injects a counted current-reference signal into the priority dynamics, rather than asserting mean conservation of plain asynchronous consensus. It constructs a new delayed, constrained latch with pre-admission physical exclusion, rather than claiming that the two supplied addenda can be used together unchanged. Finally, it reuses bodies in permanently assigned disjoint cells, rather than claiming arbitrary entry from free space. These changes are premises of the result, not consequences inherited from the supplied papers. The result is a restricted infinite-execution case under a trusted bounded coordination substrate and reserved workspace. The supplied current reference, coherent tables, commit clock, lifecycle sequencing and deadline service are its main information and timing resources. Removing or distributing components of that substrate requires subsequent results; no such implementation is derived here.

Closest external results inspected. Franceschelli and Frasca [6] already treat indefinite arrivals and departures through contractive open dynamics, bounded movement of a reference trajectory, and bounded joining error. Their open stability framework supplies a direct comparator for the flow-plus-event analysis here. The present proof uses an unnormalized energy on a finite body pool, so deleting coordinates is nonexpansive without a lower population-ratio premise. This does not make the general principle of open stability new.

Hespanha, Liberzon, and Teel [7] establish input-to-state stability conditions for impulsive systems and the role of dwell-time restrictions when jumps counteract flow decay. Here removal and exact current-latch minimization are nonexpansive in a particular energy. Admission and goal changes have explicit additive costs, and the delayed latch contributes a quantified perturbation. The scalar comparison argument is elementary and is not claimed as a new ISS theorem.

Şekercioglu, De Carli, and Dimarogonas [8] study open double-integrator formations with collision and connectivity constraints, a formation manager, prospective edges, and transition-dependent average dwell time. Their geometry is less restrictive than the reserved-cell construction used here. Our construction instead makes raw history expiry, finite label reuse, a delayed common latch, and a reference-assisted priority layer explicit. Neither paper’s guarantees are substituted for the other’s. The initial draft lists Gong, Xia, and Lv [9] as a further formation/disturbance comparator and reports metadata/supplement inspection. That supplement is not furnished here and its reading was not independently repeated in this revision; no original-body comparison or theorem dependency is asserted for that citation.

Sequence reuse after old traffic has drained is also established protocol practice. RFC 9293 discusses old connection incarnations and quiet time, while RFC 1982 specifies finite serial-number arithmetic [10, 11]. The finite-label lemma below is a proof for the explicitly bounded channel in this paper, not a claim to have invented quarantine.

Claim being made. The object studied here is the specified composition: delayed reference tracking, exclusion from physical centering, a constrained latch, an all-body invariant, and repeated use of bounded operational identities. The proofs identify the assumptions used by each conclusion. No first-publication or comparative superiority claim is needed for these bounds.

2 Execution model and the quantities to be controlled

There are $N \geq 2$ physical bodies indexed by a fixed set $\mathcal{B} = \{1, \dots, N\}$. Every body exists for all $t \geq 0$. Its physical position and velocity are $x_i, v_i \in \mathbb{R}^3$. A body may successively implement different logical incarnations, but birth and retirement do not reset x_i or v_i .

Physical commit times are $t_k = kh$, with $h > 0$. Between commits the physically admitted set $\mathcal{C}_k$ is fixed and nonempty. The logically participating set $\mathcal{L}_k$ satisfies

$$\emptyset \neq \mathcal{C}_k \subseteq \mathcal{L}_k \subseteq \mathcal{B}. \quad (1)$$

A logical newcomer may communicate and acquire priorities before entering $\mathcal{C}_k$. Bodies outside $\mathcal{C}_k$ follow a declared standby or approach controller. At most N bodies can enter or leave at one commit, and there are finitely many primitives per sampling cell. Hence infinitely many events cannot accumulate at a finite physical time.

The sets in (1) are paired epoch rosters: $\mathcal{L}_k$ is the logical computation roster established at source stage σ_k , and $\mathcal{C}_k$ is the physical command roster installed at t_k . Their inclusion is an index-matched condition. During $[\sigma_k, t_k]$ a logically closed body may still finish its previously held physical command in $\mathcal{C}_{k-1}$. It generates no new logical row after close. This transition does not assert continuous inclusion of the old physical roster in the newly selected logical roster; every body nevertheless remains in the all-body safety model.

The fixed objective alphabet has integer size $m \geq 1$. Here $\Delta_m = \{w \in \mathbb{R}^{1 \times m} : w \geq 0, w \mathbf{1}_m = 1\}$. A probability row $w_i(k) \in \Delta_m$ supplies a commanded goal

$$q_{i,k} = w_i(k)Z_i, \quad Z_i \in \mathbb{R}^{m \times 3}, \quad \max_{i,a} \|(Z_i)_{a,:}\| \leq Q. \quad (2)$$

The slot offset $r_i \in \mathbb{R}^3$ and profile Z_i belong to the physical role, not to an ever-growing incarnation index. The closed simplex is used; strict positivity can be imposed on inputs and mixing without changing the norm arguments. Birth is not a new objective coordinate in this model.

A shared translation latch s_k belongs to a nonempty closed convex set $\mathcal{S} \subseteq \mathbb{R}^3$. For the finite-workspace theorem $\mathcal{S}$ is compact and

$$\sup_{s \in \mathcal{S}} \|s\| \leq S. \quad (3)$$

For each $i \in \mathcal{C}_k$, the plant on $[t_k, t_{k+1})$ is

$$\dot{x}_i = v_i, \quad \dot{v}_i = K(r_i + q_{i,k} + s_k - x_i) - Cv_i + d_i(t), \quad K, C > 0. \quad (4)$$

Goals and latch are held on that cell. Disturbances are locally integrable. The physical safety construction later specializes to $K = ab$, $C = a + b$, $a, b > 0$, and bounded disturbances. The principal tracking inequality itself needs only $K, C > 0$.

Definition 1 (Current command error and physical event cost). Immediately after commit k , write

$$z_{i,k} = x_i(t_k) - r_i - q_{i,k} - s_k, \quad E_k^2 = \sum_{i \in \mathcal{C}_k} (K \|z_{i,k}\|^2 + \|v_i(t_k)\|^2). \quad (5)$$

This is an unnormalized aggregate error. It does not treat a departing body's disappearance from the performance index as disappearance from physical space. The all-body safety statement uses every element of $\mathcal{B}$.

Information and commit resources. A correct service establishes the finite current role table and a common commit decision. Before commit t_k , it receives complete same-time physical snapshots at an actual source time $\sigma_k \in [t_k - H, t_k]$, where $0 \leq H < h$. A frame contains the physical slot, live incarnation when applicable, intended commit identity, $x_i(\sigma_k)$, $v_i(\sigma_k)$, and the planned goal/role metadata. The planned $q_{i,k}$ is known before commitment and is used consistently by the service and the body. It need not have been applied by the plant at σ_k .

One epoch and its ordered stages. For $k \geq 1$, logical update epoch k is assigned to physical commit $t_k = kh$. Its source stage is σ_k ; these stages are strictly ordered because $H < h$. First the service applies the predecessor-bound logical close/birth and neutral-row changes selected for epoch k , forming $\mathcal{L}_k$ and ν_k . It then computes each continuing $w_i(k)$ from immutable earlier rows, initializes newborn rows from an available reference, and fixes the planned $q_{i,k} = w_i(k)Z_i$. At that same physical source time it samples all required x_i, v_i under the old held physical commands. A causal admission/guard decision uses only this bound table, the planned goals, and the old latch. Delivery, validation, decision computation, and dissemination finish by t_k . The old commands continue throughout the intervening flow. At t_k , finish that flow, apply the already selected physical removals/admissions, install the planned goals, and hold or recompute the latch. No priority, goal, incarnation, or admission-table version is changed between its source certificate and this commit.

Epoch indices also label raw production/expiry ticks; a physical transport primitive in $(t_{k-1}, t_k]$ has the ordered epoch label k , and initial primitives at t_0 have label zero. Source-row index r names the row's logical generation epoch, not the later epoch of a retransmission. Counter ordering and age comparisons use these labels. The initial commit $k = 0$ is supplied coherent data at $t_0 = 0$: $s_0 \in \mathcal{S}$, the initial rosters/goals, and the cell-admissible $x_i(0), v_i(0)$. No negative-time physical observation is presumed. Startup priority/reference registers are specified in section 5. Missing, late, conflicting, or differently bound tables cannot be committed; a hold/rejection policy then defines a different availability class unless it separately satisfies the assumptions of the claimed estimate and liveness bound.

One realization computes the logical update and planned goals at the source stage, collects the table, and disseminates the decision within the remaining H seconds. All local plants continue their old held commands until t_k . This is a bounded processing and communication contract with a common clock, not instantaneous global sensing. The paper does not derive it from radio range or from connectivity of an unconsumed graph. Directional delays are allowed inside the contract. Competing epoch frames cannot be combined into one table. A newly admitted incarnation must already be identified at the snapshot source stage. No late, incomplete, or differently bound table is applied.

Task reference. A correct reference service retains one current neutral probability row $a_i(k)$ for each $i \in \mathcal{L}_k$ and forms

$$\nu_k = \frac{1}{|\mathcal{L}_k|} \sum_{i \in \mathcal{L}_k} a_i(k). \quad (6)$$

Other explicitly supplied probability references could be used. Here birth adds a current row, retirement removes one, and an authorized current update replaces one. The service does not answer arbitrary queries about erased earlier neutral rows. Reference changes are real inputs to the tracking problem.

3 Two obstructions that determine the theorem's scope

Proposition 1 (Fresh-information and actuation obstruction). *Consider two admissible tasks whose desired position immediately after an event is respectively Ru and $-Ru$, where $R > 0$ and $\|u\| = 1$. Suppose their observations are identical for a latency interval of length $\delta > 0$. Every deterministic causal controller has worst-case position error at least R at each time strictly before disclosure. For a body starting at rest at the origin, even immediate disclosure cannot reduce the endpoint error at time t below*

$$\max\{0, R - \min(Vt, At^2/2)\} \quad (7)$$

under speed and acceleration bounds $V, A > 0$.

Proof. The controller produces the same position x in both hidden executions. The triangle inequality gives $2R \leq \|x - Ru\| + \|x + Ru\|$, so at least one error is at least R . Starting from rest, integration of

$\|\dot{v}\| \leq A$ gives $\|x(t)\| \leq At^2/2$; the speed bound gives $\|x(t)\| \leq Vt$. Reverse triangle inequality proves (7). These are lower bounds, not an asserted optimal rest-to-rest trajectory formula. $\square$

Thus an arbitrarily small uniform error guarantee cannot ignore target jump size, available time, or information age. Repeated events can maintain a nonzero tracking error even though all states remain bounded.

Proposition 2 (A free repeatedly recomputed latch can drift without bound). *For one body, $r = q = 0$, $K = 1$, $C = 2$, $h = 1$, and no constraint on s , recompute $s_k = x(t_k)$ at every sample. A constant scalar force $d > 0$ produces unbounded common position, while all variables relative to the newly set latch remain bounded.*

Proof. Write $v_k = v(t_k)$ and $\Delta x_k = x(t_{k+1}) - x(t_k)$. Exact flow from zero initial error to the new latch gives

$$v_{k+1} = de^{-1}, \quad \Delta x_k = e^{-1}v_k + d(1 - 2e^{-1}). \quad (8)$$

After the first step, $\Delta x_k = d(1 - e^{-1})^2 > 0$. Summing these increments proves unbounded position. The example already has zero communication delay and no membership changes. $\square$

A bound on relative shape or command error therefore does not imply finite absolute workspace. The compact set (3) is a mission constraint introduced here. It is not hidden inside the term “translation equivariance.” Translating both the environment and $\mathcal{S}$ preserves the coordinate description; holding a fixed world region while translating only the robots is a different task.

4 Operational retirement, finite labels, and retained influence

Definition 2 (Raw lifetime and operational expiry). A raw record is a packet, retransmission, pending callback, uncommitted table, or other object that can still be consumed under its original incarnation identity. Every such copy must cease to exist or become permanently ineligible within D_{raw} clock ticks of its original generation. Forwarding does not refresh generation time. Closing an incarnation at tick c stops all further production under that identity by $c + P$. Operational expiry means that no raw record of that incarnation can later be consumed. It does not mean that every numerical effect of its prior participation has vanished.

Theorem 1 (Finite-label reuse under bounded raw lifetime). *Assume correct endpoints and immutable source/destination bindings, the raw-lifetime contract above, bounded close processing P , and a quiet period of $q_{\text{quiet}} > P + D_{\text{raw}}$ ticks before a physical slot begins its next incarnation. Clear its old per-incarnation buffers and pending operations. Then no raw record from an earlier incarnation of that slot can be mistaken for a record of the new incarnation, for an infinite sequence of reuses. Two alternating incarnation colors suffice. With at most one source-epoch value per tick, a modulo counter of size*

$$M > 2(P + D_{\text{raw}} + 1) \quad (9)$$

can identify the finite actionable epoch window without an ever-growing stored counter.

Proof. Every old raw record is generated no later than $c + P$, so it is permanently ineligible by $c + P + D_{\text{raw}} < c + q_{\text{quiet}}$. This includes every forwarded or queued copy, by the definition of raw lifetime. A new incarnation therefore begins after all raw traffic from the immediately preceding one has expired; traffic from older ones has expired even earlier. Clearing local pending operations removes delayed callbacks that would otherwise escape the network argument. Induction over successive closes proves the claim for all finite prefixes, hence for every event of an infinite execution.

At an actual receive/use tick, an eligible epoch lies within a finite window shorter than $M/2$. Reducing both tick and generation number modulo M therefore preserves the unique age and its ordering in that window. Equal residues from an earlier wrap are excluded by lifetime, not by a probabilistic argument. The analysis may use an unbounded integer tick index without the controller retaining that integer. A periodic clock phase, modulo epoch, bounded timers, and the two colors implement the required discrete state. More than one frame per source/tick requires a separately bounded within-epoch index. $\square$

Remark 1 (The premise includes storage and retransmission). A delivery bound for packets that happen to be delivered is not a raw-lifetime bound. Nor does a timer that is restarted on forwarding establish the theorem. A malicious endpoint retaining an old signed message indefinitely is outside this compliant model. Byzantine agreement, cryptographic freshness, crash-recovery clock reconstruction, and arbitrary network partitions are not proved here.

Proposition 3 (Expiry is not erasure of incorporated influence). *Operational expiry does not imply that a retired participant has no causal effect on later numerical state. For example, a retained latch that was previously set using that participant remains a function of the old data after every packet expires. Likewise a scalar state updated as $y_{k+1} = \alpha y_k + (1 - \alpha)b$, $0 < \alpha < 1$, retains coefficient $\alpha^k > 0$ on its initial value at every finite k .*

Proof. Holding the latch does not modify its value. Iterating the scalar recurrence gives $y_k = \alpha^k y_0 + (1 - \alpha^k)b$. These statements do not depend on a remaining raw record. $\square$

Accordingly, this article does not use “causal death” to mean total noninfluence. It proves exclusion of stale operational identities and explicitly permits influence already incorporated into current priorities, physical state, or an authorized summary. The historical-memory lower bounds in [1, 5] are consistent with this distinction.

5 Delayed priority tracking without a final fixed roster

Let $D \geq 1$ be the maximum age in logical update ticks. For a continuing live incarnation, use

$$w_i(k) = (1 - \zeta) \sum_{(j,r) \in \mathcal{A}_{i,k}} \alpha_{i,j,r}(k) w_j(r) + \zeta \hat{\nu}_{i,k} + e_{i,k}^{\text{rnd}}, \quad 0 < \zeta \leq 1. \quad (10)$$

Coefficients are nonnegative and sum to one. Every used row has $k - D \leq r \leq k - 1$ and an eligible incarnation. The body can use its own preceding row when no remote row is available. The reference is a probability row and satisfies

$$\left\| \hat{\nu}_{i,k} - \nu_{s(i,k)} \right\|_1 \leq \epsilon_\nu, \quad k - D \leq s(i,k) \leq k. \quad (11)$$

The optional post-computation rounding preserves the simplex and has $\left\| e_{i,k}^{\text{rnd}} \right\|_1 \leq \epsilon_w$. A newborn initializes directly from an eligible reference, with the same rounding guarantee; it does not inherit the previous incarnation’s local row as an unmarked bootstrap. Every continuing live body updates at every logical tick. A schedule with less frequent updates needs its own blocking bound.

Set

$$e_k = \max_{i \in \mathcal{L}_k} \|w_i(k) - \nu_k\|_1, \quad V_k = \sum_{\ell=k-D+1}^k \|\nu_\ell - \nu_{\ell-1}\|_1, \quad \epsilon = \epsilon_\nu + \epsilon_w. \quad (12)$$

Before the first update, supply probability references $\nu_{-D}, \dots, \nu_0$ and a finite set of initial source registers indexed by $r \in \{-D, \dots, 0\}$. Each register is either an actual immutable row of an incarnation

live at its source index, with its then-reference, or a specifically declared startup row at a negative index. Every index-zero register must equal the installed $w_i(0)$ of an incarnation in $\mathcal{L}_0$; an additional synthetic index-zero value is not eligible. A constant extension $\nu_r = \nu_0$, negative-index startup rows equal to ν_0 , and current rows initialized from ν_0 is one admissible convention. Every startup row and its error bound is part of the initial input; it is not an undeclared past execution. For negative indices define e_r as the maximum startup/source-register error at that index (zero for an unused empty index); e_0 remains the actual live-row maximum in (12). Let e_{init} bound these negative-index errors and e_0 . Thus the first V_k and every initial delayed window are defined.

An actionable numerical source row must be bound to the incarnation actually live at its generation index r , its immutable value $w_j(r)$, and the reference ν_r , except for the explicitly supplied startup registers. A source removed at close cannot generate a new logical row outside $\mathcal{L}_r$ for use in (10). Close processing may send or forward copies of a previously valid immutable row until its transport deadline; it does not replace the row's source index by the copy time. Destination incarnation, current task metadata, raw lifetime, and the D -epoch source age are checked separately. A later slot reuse never reinterprets a prior row as the new incarnation's state. These conditions are what permit the proof to bound each consumed row by e_r .

Theorem 2 (Reference-assisted tracking under perpetual logical turnover). *Under (10)–(12),*

$$e_k \leq (1 - \zeta) \max_{k-D \leq r < k} e_r + V_k + \epsilon. \quad (13)$$

In particular, $V_k + \epsilon \leq \beta$ for all $k \geq 1$ gives

$$e_k \leq c + (1 - \zeta)^{\lfloor k/D \rfloor} (e_{\text{init}} - c)_+, \quad c = \beta/\zeta. \quad (14)$$

The trivial simplex bound $e_k \leq 2$ may be intersected with this estimate. Exact reference representation, zero rounding, and $V_k \rightarrow 0$ imply $e_k \rightarrow 0$, even without a final fixed live set. Also $\|w_i(k) - w_j(k)\|_1 \leq 2e_k$ for current live rows.

Proof. For any eligible source time r , telescoping yields $\|\nu_r - \nu_k\|_1 \leq V_k$. The same inequality holds for the reference time $s(i, k)$. Subtract ν_k from (10), use convexity, and add the two bounded representation errors. The total coefficient on reference-time variation is one, which gives (13). A newborn has error at most $V_k + \epsilon$, so the same upper bound applies. Deleting rows from the current maximum does not increase it.

Subtract $c = \beta/\zeta$ and take positive parts. The sequence $f_k = (e_k - c)_+$ satisfies $f_k \leq (1 - \zeta) \max_{k-D \leq r < k} f_r$. In every block of D updates all members of the actionable error window are replaced by values at most $1 - \zeta$ times the preceding block maximum. Induction proves (14). For vanishing V_k and $\epsilon_\nu = \epsilon_w = 0$, apply that estimate to an arbitrarily late tail with arbitrarily small β ; the finite preceding history decays. The pair inequality is the triangle inequality through ν_k . For $\zeta = 1$, continuing rows depend only on the reference and rounding; the same conclusion follows directly, with the usual convention for a zero power at the initial block. $\square$

A nonzero reference-representation error cannot be dropped from the recovery premise. With constant $\nu = (1/2, 1/2)$ and injected $\hat{\nu} = (3/5, 2/5)$, a self-mixing row converges to $\hat{\nu}$, with limiting error $1/5$, even with exact arithmetic and no reference variation. Exact recovery therefore uses $\epsilon_\nu = \epsilon_w = 0$; time-varying bounds tending to zero give the same conclusion by the tail argument.

The service (6) is a counted informational resource. No connectivity hypothesis is needed for theorem 2 because every live node receives reference injection, not because disconnected plain consensus has acquired a new property. Dynamic objective supports and proportional conditioning would require different transport/reset estimates. They are not silently absorbed into a fixed-simplex theorem.

For an entrant with neutral row a and n prior live rows, one has

$$\nu^+ - \nu^- = (a - \nu^-)/(n + 1). \quad (15)$$

For a departure from $n > 1$ live rows,

$$\nu^+ - \nu^- = (\nu^- - a)/(n - 1). \quad (16)$$

These are actual event-size identities, obtained by expanding the two means. Several operations are composed in their declared order. They show why event count alone does not determine reference variation.

6 A delayed constrained latch and pre-admission physical exclusion

At t_k , first complete the old plant flow. Determine departures and proposed physical admissions from the bound source table, install the new goals, and then update the common latch. No position or velocity is reset. Let

$$\mathcal{R}_k = \mathcal{C}_{k-1} \cap \mathcal{C}_k, \quad \mathcal{J}_k = \mathcal{C}_k \setminus \mathcal{C}_{k-1}, \quad n_k = |\mathcal{C}_k|. \quad (17)$$

The old latch s_{k-1} is denoted s_k^- at this event. A hold sets $s_k = s_k^-$. A recomputation uses

$$s_k = \text{proj}_{\mathcal{S}} \left(\frac{1}{n_k} \sum_{i \in \mathcal{C}_k} [x_i(\sigma_k) - r_i - q_{i,k}] \right). \quad (18)$$

A separately quantified implementation residual may replace this by any $s_k \in \mathcal{S}$ within distance δ_k of the displayed projection. Let $b_k \in \{0, 1\}$ record recomputation. Nonadmitted physical bodies are excluded from this mean. They may nonetheless affect the logical reference through declared neutral preferences. We call this pre-admission physical exclusion. It is not the earlier common-mode shield of [4]: every body here already belongs to a bounded reserved cell, and logical participation may still change incumbent goals. No equivalence to that earlier live-admission law is asserted.

The abstract projection is an explicit computation-service resource. For an arbitrary closed convex set, existence and single-valuedness do not supply a finite representation, algorithm, or deadline. A concrete compact option is $\mathcal{S} = \overline{B}_S(0)$, $S > 0$, with

$$\text{proj}_{\mathcal{S}}(a) = \begin{cases} a, & \|a\| \leq S, \\ Sa/\|a\|, & \|a\| > S. \end{cases}$$

This formula gives the exact-real service used by the simple realization. With rational inputs and radius, its square root can instead be bracketed to compute an inward admissible output with a certified δ_k ; the finite precision, work, and time budget must then be included in the service contract. The general projection theorem is retained without a computability claim for every convex set.

A nonvacuous guard family. The tracking estimate permits any causal sequence of holds and recomputations. For the quiescent-tail result, require a settled-body exclusion: a recomputation is forbidden when every admitted body, in the coherent source table, has error to its planned goal plus the old latch at most $R_g > 0$ and speed at most $V_g > 0$. Otherwise either branch is allowed. The radial and timer rules in [3] are a stricter possible policy, but this article does not claim exact callback correspondence. Near a guard threshold, an uncertain measurement requires a stated robust decision rule.

Lemma 1 (Uniform exact-flow contraction). *Let*

$$A = \begin{bmatrix} 0 & 1 \\ -K & -C \end{bmatrix}, \quad P = \text{diag}(K, 1), \quad \sigma_h = \|P^{1/2} e^{Ah} P^{-1/2}\|_2. \quad (19)$$

For $K, C, h > 0$, $\sigma_h < 1$. The value does not depend on cohort size. For a held goal and latch over one cell,

$$E(t_k^-) \leq \sigma_h E_{k-1} + U_{k-1}, \quad U_{k-1} = \int_{t_{k-1}}^{t_k} \left\| \text{col}_{i \in \mathcal{C}_{k-1}} d_i(t) \right\| dt. \quad (20)$$

At any intermediate point in that cell, the corresponding error norm is at most $E_{k-1} + U_{k-1}$.

Proof. For a homogeneous scalar coordinate, differentiation gives $\frac{d}{dt}(Kz^2 + v^2) = -2Cv^2$. Thus flow is nonexpansive. Equality over a positive interval would force $v = 0$ throughout, then $\dot{v} = -Kz = 0$, hence the entire state is zero. The continuous norm ratio therefore has maximum strictly below one on the compact P -unit sphere. Repeating the same two-dimensional block for three axes and any number of bodies leaves its operator norm unchanged. The forcing injection $d \mapsto (0, d)$ has energy norm one. Variation of constants, nonexpansion for intermediate flow, and the triangle inequality give both claims. $\square$

Lemma 2 (Current projection and delayed projection defect). *For any nonempty cohort, fixed x_i, q_i, r_i, v_i , and old $s \in \mathcal{S}$, replacing s by $\text{proj}_{\mathcal{S}}(n^{-1} \sum_i (x_i - r_i - q_i))$ cannot increase the error energy. Using a complete older position table at $\sigma \leq t$, but the same planned q_i , adds at most*

$$\sqrt{K} \int_{\sigma}^t \left\| \text{col}_i v_i(u) \right\| du + \sqrt{nK} \delta \quad (21)$$

to the norm, where δ bounds the projection implementation residual in Euclidean norm.

Proof. Put $a_i = x_i - r_i - q_i$ and $\bar{a} = n^{-1} \sum_i a_i$. The identity $\sum_i \|a_i - s\|^2 = \sum_i \|a_i - \bar{a}\|^2 + n \|\bar{a} - s\|^2$ reduces minimization to Euclidean projection on $\mathcal{S}$. Velocities are unchanged. Projection onto a closed convex set is nonexpansive: its variational inequalities at projections p, p' of a, a' give $\|p - p'\|^2 \leq (p - p')^\top (a - a') \leq \|p - p'\| \|a - a'\|$. The current and old means differ by at most $n^{-1/2} \int_{\sigma}^t \left\| \text{col}_i v_i(u) \right\| du$. Repeating a latch difference in n position rows has energy norm $\sqrt{nK}$ times that difference. Combine nonexpansion, the exact minimization property, and the implementation residual. $\square$

Define the event costs, using the actual state at t_k before the latch reset,

$$P_k = \sqrt{K} \left\| \text{col}_{i \in \mathcal{R}_k} (q_{i,k} - q_{i,k-1}) \right\|, \quad (22)$$

$$A_k = \left(\sum_{j \in \mathcal{J}_k} \left[K \left\| x_j(t_k) - r_j - q_{j,k} - s_k^- \right\|^2 + \|v_j(t_k)\|^2 \right] \right)^{1/2}, \quad (23)$$

$$W_k \geq \sup_{u \in [\sigma_k, t_k]} \left\| \text{col}_{j \in \mathcal{J}_k} v_j(u) \right\|. \quad (24)$$

Empty entrant sets have $A_k = W_k = 0$. A certificate may overestimate these quantities, but it cannot omit one merely because it is inconvenient to observe.

Theorem 3 (Physical exclusion and a uniform latch event inequality). *Under the stated flow, source coherence, event order, and projection contracts, set $d_H = \sqrt{K}H$. Every execution satisfies*

$$E_k \leq (\sigma_h + b_k d_H) E_{k-1} + J_k, \quad (25)$$

$$J_k = (1 + b_k d_H) U_{k-1} + P_k + A_k + b_k d_H W_k + b_k \sqrt{n_k K} \delta_k. \quad (26)$$

In particular, $\rho = \sigma_h + d_H < 1$ implies

$$E_k \leq \rho^k E_0 + \sum_{\ell=1}^k \rho^{k-\ell} J_\ell. \quad (27)$$

The bound is independent of the number of prior incarnations. Before physical admission, an excluded body's absolute position and velocity do not occur in incumbent feedback or in the latch mean, conditional on the actual goal, disturbance, commit, and branch signals.

Proof. At the end of the old cell, lemma 1 bounds the old-cohort error by $\sigma_h E_{k-1} + U_{k-1}$. Deleting departed rows cannot increase this unnormalized norm. Changing the goals of the continuing rows adds at most P_k . Appending the entrant rows changes a nonnegative norm e to $\sqrt{e^2 + A_k^2} \leq e + A_k$. Thus the post-membership, pre-latch error is at most $\sigma_h E_{k-1} + U_{k-1} + P_k + A_k$.

For a recomputation, use lemma 2. Since $H < h$, every continuing body was in the preceding physical cohort throughout $[\sigma_k, t_k]$. Its stacked speed is bounded by $E_{k-1} + U_{k-1}$ through lemma 1. The entrant contribution is at most W_k . The triangle inequality therefore bounds the old-to-current projection defect by $d_H(E_{k-1} + U_{k-1} + W_k) + \sqrt{n_k K} \delta_k$. Adding gives (25). A hold has no such term. Iteration with $\sigma_h + b_k d_H \leq \rho$ proves (27). The final claim follows directly from (4) and (18): excluded physical coordinates occur in neither expression for an incumbent. It is not a claim against indirect influence through a motion-dependent communication schedule or a goal inferred from that body's position. $\square$

The factor $\rho < 1$ is sufficient, not necessary. A run with $\rho \geq 1$ is uncertified by (27); this inequality alone does not prove that run unstable. Complete old tables, current local feedback, and the event order are all used. Mixing different source epochs or retaining a departed body's row in the latch is a different controller.

6.1 Event size instead of event count

Corollary 1 (Windowed event-cost budget). *Suppose $\rho < 1$ and $J_k \leq u + \eta_k$, where $u \geq 0$, $\eta_k \geq 0$, and every L -tick window satisfies*

$$\sum_{j=\max(1, k-L+1)}^k \eta_j \leq B_{\text{evt}}. \quad (28)$$

Then

$$E_k \leq \rho^k E_0 + \frac{u}{1-\rho} + \frac{B_{\text{evt}}}{1-\rho^L}. \quad (29)$$

A uniformly bounded per-step J_k is sufficient without any requirement that changes eventually stop. Under $J_k \rightarrow 0$, $E_k \rightarrow 0$.

Proof. In (27), the constant input contributes at most $u/(1-\rho)$. Partition the remaining sum backward into blocks of length L . Block r , counting the most recent as zero, has total input at most B_{evt} and discount at most ρ^{rL} . Sum the resulting geometric series. For a vanishing input, split the convolution into a finite prefix, which decays, and a tail bounded by an arbitrarily small constant times $1/(1-\rho)$. $\square$

The event part of (29) is attained by the scalar comparison recursion at its peaks when an input B_{evt} occurs every L steps. It is not asserted to be attainable by every physical formation. An average event count becomes a useful proxy only after a bound on the size of each corresponding P_k, A_k, W_k has been supplied. The proposed article's earlier dwell-time sketch is therefore replaced by a concrete additive event budget in this construction.

7 All-body safety in a finite reusable workspace

The performance proof removes retired rows from E_k ; the safety proof must not remove their bodies. Assume $K = ab$, $C = a + b$, $a, b > 0$. Every physical body, including standby and approach bodies, uses

$$\dot{x}_i = v_i, \quad \dot{v}_i = ab(c_i(t) - x_i) - (a + b)v_i + d_i(t), \quad \|d_i(t)\| \leq D_b. \quad (30)$$

All possible held or measurable goals satisfy $\|c_i(t) - r_i\| \leq B$, with $B \geq Q + S$ for admitted bodies. Standby goals satisfy the same envelope. Choose

$$R \geq B + D_b/(ab), \quad C_i = r_i + \bar{B}_R(0), \quad y_i = x_i + v_i/b, \quad (31)$$

and require $x_i(0), y_i(0) \in C_i$. Each body has radius $r_b > 0$. Slot centers obey

$$\|r_i - r_j\| \geq 2R + 2r_b + \delta_{\text{sep}}, \quad \delta_{\text{sep}} > 0. \quad (32)$$

There are no other bodies or obstacles in the reserved cells.

Theorem 4 (Robust safety through arbitrarily many mode changes). *Under (30)–(32), every physical body remains in its fixed cell for all time, and*

$$\|v_i(t)\| \leq V_{\text{cert}} := 2bR, \quad (33)$$

$$\|u_i(t)\| \leq U_{\text{cert}} := ab(B + R) + 2b^2R, \quad (34)$$

$$\|\dot{v}_i(t)\| \leq A_{\text{cert}} := 2b(a + b)R \quad a.e. \quad (35)$$

Here $u_i = ab(c_i - x_i) - (a + b)v_i$ is commanded acceleration. All-body center separation is at least $2r_b + \delta_{\text{sep}}$. These statements hold independently of the number of retirements, births, or latch changes. They do not require $\rho < 1$.

Proof. Transform (30) to

$$\dot{y}_i = a(\tilde{c}_i - y_i), \quad \dot{x}_i = b(y_i - x_i), \quad \tilde{c}_i = c_i + d_i/(ab). \quad (36)$$

The radius condition puts $\tilde{c}_i(t)$ in C_i almost everywhere. Variation of constants expresses $y_i(t)$ as a convex average of its initial value and previous values of $\tilde{c}_i$. It therefore stays in the closed convex cell. Apply the same argument to x_i , with input y_i . This proof covers every allowed change in c_i and bounded measurable disturbance, not just samples.

Since $v_i = b(y_i - x_i)$, (33) follows. Algebra gives $u_i = ab(c_i - y_i) + b^2(x_i - y_i)$ and $\dot{v}_i = ab(\tilde{c}_i - y_i) + b^2(x_i - y_i)$. The first distance is at most $B + R$ for the command, and at most $2R$ for the effective input. The other distance is at most $2R$. These estimates give (34)–(35). Finally, $\|x_i - x_j\| \geq \|r_i - r_j\| - 2R$ proves separation. The union of finitely many bounded cells is a fixed bounded workspace. No physical state is reset at a mode switch. $\square$

Actuators with command capacity at least U_{cert} never saturate on this nominal model. A specified physical acceleration limit must exceed A_{cert} . These are sufficient design inequalities, not a treatment of thrust allocation, attitude, gravity estimation, aerodynamic coupling, contacts, or battery depletion. Infinite execution presumes the ideal plant and stated energy/service availability continue to operate. A finite physical workspace is not a claim of infinite energy from finite batteries.

The invariance premise on y_i is substantive. A body at the boundary of its cell with outward velocity cannot remain there under any bounded acceleration merely because its position initially belongs to the cell. The safety region therefore constrains both position and velocity.

8 An executable gate and a finite-service recycling construction

The stale physical table also contains velocity. The bounds in theorem 4 are available before admission and do not depend on the tracking certificate. For a proposed entrant j , define from its source record

$$\widehat{z}_{j,k} = x_j(\sigma_k) - r_j - q_{j,k} - s_k^-, \quad (37)$$

$$\mathbf{a}_{j,k} = \left[K(\|\widehat{z}_{j,k}\| + V_{\text{cert}}H)^2 + (\|v_j(\sigma_k)\| + A_{\text{cert}}H)^2 \right]^{1/2}, \quad (38)$$

$$\mathbf{w}_{j,k} = \|v_j(\sigma_k)\| + A_{\text{cert}}H. \quad (39)$$

A numerical implementation needs outward rounding of these certificates; the exact-arithmetic theorem does not assume a floating-point comparison is automatically conservative.

Theorem 5 (Delayed-observation admission certificate). *An admission selected at the source stage and committed at t_k satisfies*

$$A_k \leq \left(\sum_{j \in \mathcal{J}_k} \mathbf{a}_{j,k}^2 \right)^{1/2}, \quad W_k \leq \left(\sum_{j \in \mathcal{J}_k} \mathbf{w}_{j,k}^2 \right)^{1/2}. \quad (40)$$

Thus the service can certify specified admission-cost limits without an instantaneous centroid oracle. It must apply only the declared source table and flush or reject all older admission-table versions. Expired or mismatched incarnation records cannot be used to pass the gate.

Proof. Over the source-to-commit interval, integration gives $\|x_j(t_k) - x_j(\sigma_k)\| \leq V_{\text{cert}}H$ and $\|v_j(t_k) - v_j(\sigma_k)\| \leq A_{\text{cert}}H$. The old latch is fixed over that interval because $H < h$; the planned goal is fixed metadata. Applying these two bounds in (23) proves the first claim. The same acceleration estimate holds at every point between source and commit, proving the second. Version and incarnation checks ensure these inequalities are attached to the state and goal that will actually be admitted, rather than a different pending update. $\square$

Feasibility versus precision. A gate with an arbitrarily small threshold need not ever pass under persistent forcing or changing goals. This article does not infer liveness from convergence under a hypothetical quiet tail. The following explicit broad threshold is always feasible within the reserved cells:

$$a_{\text{broad}} = \left[K(R + B + V_{\text{cert}}H)^2 + (V_{\text{cert}} + A_{\text{cert}}H)^2 \right]^{1/2}. \quad (41)$$

Indeed $\|\widehat{z}_{j,k}\| \leq R + B$ and source speed is at most V_{cert} . A one-body gate at a_{broad} therefore always passes. A smaller threshold trades a tighter event budget for a separate capture-time obligation. It may be used in a particular certified schedule, but is not granted unlimited-time liveness for free.

8.1 Repeated physical reuse

For this canonical recycling construction, all N bodies are initially physically admitted. A correct reservation service accepts at most one pending lifecycle request per body and processes requests in FIFO order, one at a time. It never removes the final admitted body. For a served body it selects physical exclusion and closes the old logical incarnation at a source stage; at the ensuing commit it removes that body from the physical mean and commands a fixed standby goal in its cell. It waits the quiet period and an optional finite standby dwell, initializes the new color from the reference service, collects the fresh planned-goal/physical table, and applies the broad gate before readmission. Initialization has a declared bound B_{boot} in clock ticks. Closing and reuse do not erase the body's physical state.

A standby dwell can represent entry into a tolerance region, not exact finite-time rest. To make that statement explicit, fix standby goal c_i^0 and let $L_i = R + \|c_i^0 - r_i\|$. For constant standby goal, (36) gives

$$\|y_i(t) - c_i^0\| \leq L_i e^{-at} + D_b/(ab), \quad (42)$$

$$\|x_i(t) - c_i^0\| \leq L_i \phi(t) + D_b/(ab), \quad (43)$$

$$\|v_i(t)\| \leq bL_i(e^{-at} + \phi(t)) + 2D_b/a, \quad (44)$$

$$\phi(t) = e^{-bt} + b \int_0^t e^{-b(t-s)} e^{-as} ds \longrightarrow 0.$$

These estimates follow by the same convex-filter convolution and $v_i = b(y_i - x_i)$. Therefore tolerances strictly larger than $D_b/(ab)$ in position and $2D_b/a$ in speed admit a finite common dwell for the finite body set. Stronger damping estimates can improve these sufficient margins.

Theorem 6 (Bounded lifecycle service and finite resource reuse). *Assume the channel/initialization contracts, the all-body cell premises, an initially fully admitted physical pool, the serialized FIFO reservation service, a finite chosen standby dwell, and the broad admission gate. Each accepted request finishes in a bounded finite number of ticks. With at most one pending request per body, the waiting plus service time is bounded independently of execution length. An infinite sequence of such accepted requests can be completed by the same N physical bodies in the same fixed workspace. No new berth or body is consumed by a completed request.*

Proof. Each listed service phase has a finite fixed bound: next-commit alignment costs at most one tick; close processing and quarantine are covered by the fixed quiet period; standby dwell is fixed; bootstrap costs at most B_{boot} ; collecting and committing the fresh table costs at most two further tick alignments under $H < h$. For example

$$L_{\text{srv}} = q_{\text{quiet}} + \lceil T_{\text{standby}}/h \rceil + B_{\text{boot}} + 4 \quad (45)$$

is a conservative service bound. The broad gate passes by (41). Every completed cycle restores all N bodies to the admitted core. During a cycle only its served body is excluded. Consequently $N \geq 2$ preserves at least one admitted body. The initially fully admitted premise is used here; mere nonemptiness would not exclude a last-body deadlock.

There are at most N accepted outstanding requests. Fair FIFO service therefore finishes any accepted request within NL_{srv} ticks. A new request can be accepted when its body's prior request finishes; repeatedly cycling through the body set gives an explicit infinite accepted sequence. Theorem 1 supplies safe logical reinitialization at each cycle. Theorem 4 applies through every physical phase and retains every body. No position, velocity, slot, or reserved region is recreated. The proof does not serve an unbounded external request queue with finite memory: requests beyond the declared capacity are explicitly rejected or deferred outside the accepted queue. $\square$

This is resource closure under a supplied finite service discipline: the finite phase, queue and broad-gate premises supply the completion bound, while repeated cycles consume no additional body, berth or operational identity fields. It proves repeated deactivate/reactivate service, with optional standby-tube arrival. It does not prove battery replacement, exact docking, repair of failed hardware, exchanges between different physical cells, or perpetual capacity under arbitrary adversarial requests. Those are different lifecycle contracts.

9 Task-space guarantees and an infinite-execution composition

The physically relevant desired profile at commit k is $r_i + \nu_k Z_i$, up to the shared latch. For $t \in [t_k, t_{k+1})$, let $E(t)$ be the energy relative to the held command. Then

$$E(t) \leq E_k + U_k. \quad (46)$$

The U_k on the right may be replaced by its partial-cell integral for a causal certificate.

Lemma 3 (Separate relative and common task errors). *For $i, j \in \mathcal{C}_k$ and $t \in [t_k, t_{k+1})$,*

$$\|(x_i - x_j) - (r_i - r_j + \nu_k(Z_i - Z_j))\| \leq \sqrt{2/K} E(t) + 2Qe_k, \quad (47)$$

$$\|\bar{x}_{\mathcal{C}_k} - \bar{r}_{\mathcal{C}_k} - s_k - \nu_k \bar{Z}_{\mathcal{C}_k}\| \leq E(t)/\sqrt{n_k K} + Qe_k. \quad (48)$$

The common error relative to the same task with zero translation is at most the right side of (48) plus S .

Proof. The profile radius gives $\|(w_i - \nu_k)Z_i\| \leq Q \|w_i - \nu_k\|_1$. Subtract the two actual command-error identities; the common latch cancels, and $\|z_i - z_j\| \leq \sqrt{2} (\sum_{\mathcal{C}} \|z_\ell\|^2)^{1/2}$. This proves (47). Average those identities and use Cauchy–Schwarz to bound $\|\bar{z}\| \leq E(t)/\sqrt{n_k K}$. Finally add $\|s_k\| \leq S$ for the zero-translation reference. $\square$

With shared profiles $Z_i = Z$, the task-dependent term in pair geometry vanishes exactly. With distinct fixed profiles it represents actual task-dependent deformation. Neither observation establishes equality of two different historical priority endpoints.

Theorem 7 (Persistent operation with finite reusable bodies). *For every supplied infinite execution satisfying all of the following contracts:*

1. *a finite physical pool, a nonempty admitted core, the fixed objective alphabet and bounded profiles;*
2. *correct bounded-age reference and priority service as in theorem 2;*
3. *complete coherent physical tables, the event order and constrained latch of theorem 3, with $\rho < 1$;*
4. *the robust cell, initial-state, actuator-capacity and workspace premises of theorem 4;*
5. *bounded raw lifetime, bounded close/bootstrap service, finite-label quarantine, and the initially fully admitted, capacity-bounded lifecycle scheduler of theorem 6;*
6. *declared bounds on rounding and projection residuals.*
7. *for uniform operational storage, a fixed aggregate copy-production cap F per epoch (including re-transmissions), bounded simultaneous bursts, bounded frame/within-epoch fields, and bounded serviced queues as specified in section 11; computation and dissemination meet their declared deadlines under that actual workload.*

Then every body exists for all time; position and velocity are continuous; all-body collision separation, speed and acceleration bounds hold; every accepted lifecycle request completes; and a new incarnation cannot consume raw records left over from an earlier one after the stated quarantine. Priority disagreement and physical task errors satisfy (14), (27), and (47)–(48). These conclusions require no final fixed-membership epoch.

All quantities in the uniform physical event bound can be bounded using finite resources. In the exact-projection case,

$$U_k \leq h\sqrt{N}D_b, \quad (49)$$

$$P_k \leq 2\sqrt{N}KQ, \quad (50)$$

$$A_k \leq \sqrt{N} [K(R+B)^2 + V_{\text{cert}}^2]^{1/2}, \quad (51)$$

$$W_k \leq \sqrt{N}V_{\text{cert}}. \quad (52)$$

Consequently J_k has a finite execution-length-independent bound. The event-sensitive certificate (27) and budget (29) can be much smaller than this worst-case envelope. Finite operational storage is specified in section 11; exact real variables are not claimed to have finite bit cost.

Proof. Fix any supplied execution in the stated class, including its disturbance signals, bound tables, ordered events and held commands. On each cell, the linear plant has an absolutely continuous solution for those locally integrable inputs. Concatenating the cell solutions preserves position and velocity at commits; there are only finitely many commits in any finite-time interval. This is a pathwise statement for each compliant execution, not an existence or uniqueness theorem for a coordination or guard policy. The cell invariants prevent physical escape and prove the actuator and separation claims, including for bodies omitted from E_k . Bounded lifecycle service and stale-record exclusion follow from theorems 1 and 6.

Each row and reference is a probability vector, so task variation in a finite D -tick window is finite and uniformly at most $2D$ even without a tighter signal contract. Theorem 2 therefore applies at every logical tick. The actual physical events obey theorem 3, giving the convolution; lemma 3 transfers the two estimates to the declared task.

For the explicit uniform event bounds, sum the disturbance cap over at most N bodies, use $\|q_i\| \leq Q$ for goal changes, and use $\|x_i - r_i\| \leq R$, $\|q_i + s\| \leq B$, and the certified speed for entrant states. The entrant-speed estimate is valid over the entire source interval, not just at its endpoints. Add $\sqrt{NK} \sup \delta_k$ when projection residuals are present. Thus bounded perpetual operation is not obtained by assuming the final error is bounded. Every needed cost has a prior resource or invariant bound. None of these arguments erases the distinction between admitted and nonadmitted bodies, or between a logical identity and a physical slot. $\square$

What the theorem buys, and what it does not. The system can complete endless admitted lifecycle operations and track an explicitly moving task within the stated analytical envelope. Its event-sensitive bound uses the actual discounted disturbance and event costs. A finite numerical certificate separately requires effective input representations and certified evaluation; arbitrary exact-real data or a general projection oracle is not automatically a finite algorithm. Safety follows from a continuous invariant. The price is an explicit reserved physical layout, a bounded common-translation region, and reference/commit infrastructure. These restrictions remain in the abstract and the main theorem. In particular, bounded accepted-request completion is conditional on the supplied service discipline; its resource conclusion is that indefinite repeated reuse does not increase the required physical fleet or workspace.

10 Recovery on a quiescent tail

Theorem 8 (Eventual latch freeze and nominal task recovery). *Assume the full tracking/commit frame of theorem 3 with $\rho = \sigma_h + \sqrt{K}H < 1$. Suppose physical membership becomes fixed after a finite tick, $q_{i,k} \rightarrow q_{i,\infty}$, $U_k \rightarrow 0$, and $\delta_k \rightarrow 0$. The latch guard has the settled-body exclusion described in section 6, using exactly the planned goals installed at the ensuing commit. Then there are only finitely many successful recomputations. The retained latch becomes a finite constant $s_\infty \in \mathcal{S}$ and*

$$x_i(t) \longrightarrow r_i + q_{i,\infty} + s_\infty, \quad v_i(t) \longrightarrow 0. \quad (53)$$

For exact reference representation, zero priority rounding, and a convergent neutral reference with $V_k \rightarrow 0$, the reference-assisted layer gives $q_{i,\infty} = \nu_\infty Z_i$ for each body in the final physical core.

Proof. On the fixed physical tail, $A_k = W_k = 0$. Goal convergence gives $P_k \rightarrow 0$. Hence $J_k \rightarrow 0$, so corollary 1 gives $E_k \rightarrow 0$, and (46) gives the same conclusion between samples. Speeds at source times vanish. The source error measured with the new planned goal differs from the old held-command error by $q_{i,k} - q_{i,k-1}$; that difference also vanishes. Eventually every source body is strictly inside the two positive settled thresholds. The guard then forbids any further recomputation. Only finitely many earlier samples exist, so the actual last latch is a finite constant. Combining $E(t) \rightarrow 0$ with $q_{i,k} \rightarrow q_{i,\infty}$ yields (53). The priority assertion follows from theorem 2 and the fixed linear profiles. $\square$

Vanishing increments without convergence of the goals give a tracking statement, not a position endpoint. Bounded nonzero disturbance generally leaves a nonzero error tube. Without the settled-body guard, infinitely many tiny recomputations can continue; $E_k \rightarrow 0$ alone does not imply convergence of the latch. The endpoint is the actual frozen latch plus the chosen current reference, not the plain delayed-consensus endpoint from a different update law.

11 Finite operational memory and representation limits

The reference service needs at most Nm current neutral numbers and current role metadata. The priority protocol retains at most a configured finite number of source rows per physical slot within its D -tick actionable window, for example $O(NmD)$ numbers network-wide in a central ring representation, or the corresponding replicated amount for per-node buffers. A complete physical source table contains $O(N)$ three-vectors and fixed-size role fields. At most a committed and a pending table are needed under the stated one-commit-per-cell contract. A nondeleting audit recorder is optional external storage; it is not part of this operational-state bound.

For a channel with at most F generated copies per tick, counting retransmissions as copies, raw lifetime D_{raw} bounds the number of raw in-flight/pending records by $F(D_{\text{raw}} + 1)$, up to a separately specified bounded simultaneous burst. Merely locally finite traffic would not give that uniform bound. The direct-service realization can provision F for the actual N , profile payload, reference broadcasts, membership traffic and physical snapshots; lifecycle and data traffic are both counted. No throughput claim is derived from a static communication graph.

A bounded direct-service realization. As a concrete offered-load policy, use one immutable priority row per source/destination pair per epoch (at most N^2 messages), with no retry or transport replication. Permit at most N messages in each of eight service classes per epoch: lifecycle requests, role records, acknowledgements, neutral-row updates, reference replies, physical snapshots, commit decisions, and bootstrap replies. Thus $F = N^2 + 8N$ is a sufficient production cap for this policy; smaller implementations can use their actual count. Bundled decisions contain at most the current N role/goal entries, never an unbounded historical log. Fixed dimensions, bounded per-epoch subindices, and finite slot/colour/modulo fields bound the number of fields per frame, while real-number precision remains separately counted. Overflow is rejected before admission to a finite queue. Accepted traffic is delivered or terminally discarded within the raw lifetime, and required reference/bootstrap/commit traffic is serviced within its stated budget under this load. These last correctness and scheduling guarantees are premises of the direct service, not properties inferred from its static layout or copy cap. With capped production, lifetime, and capacity-bounded requests, the construction has uniform record counts; without them, theorems 3 and 4 may still hold while raw storage peaks grow.

In this operational realization, consumed, terminally rejected, and expired records and callbacks are removed from operational buffers; clearing/disposal work is included in the service budget. Merely marking a record permanently ineligible does not bound its retained bytes. If an audit copy is kept, it belongs to the separately counted external archive and is inaccessible to operational decoding. The gen-

eral raw-lifetime theorem only requires permanent ineligibility; the stronger disposal/capacity contract is used for the uniform operational-storage conclusion.

Discrete identity fields use physical slot IDs, two incarnation colors, modulo epochs, bounded timers and the fixed-size current role table. They need not grow with completed incarnation count. Stored reference rows, latches and physical coordinates are still exact real variables in the nominal theorem. A fixed number of reals is not a fixed number of bits. A finite probability grid can implement a simplex-preserving quantizer: for $m = 2$, rounding the first coordinate to the nearest multiple of $1/L$ and setting the second to its complement gives ℓ^1 error at most $1/L$. The corresponding nonzero ϵ_u or ϵ_v is already charged in theorem 2. General simplex quantizers need a proved uniform error bound.

The physical projection must likewise keep its output inside $\mathcal{S}$ and supply its numerical error δ_k . Quantization alone does not certify a numerical integrator. With local measurement errors $\tilde{x} - x = e_x$ and $\tilde{v} - v = e_v$, replacing exact feedback by measured feedback introduces the explicit acceleration residual

$$d_{\text{eff}} = d_{\text{env}} - K e_x - C e_v + e_u. \quad (54)$$

Uniform sensor and actuator-error bounds give $\|d_{\text{eff}}\| \leq D_{\text{env}} + K \epsilon_x + C \epsilon_v + \epsilon_u$. Applying the cell theorem requires using this full bound and checking the resulting commanded-actuation margin. The existence of such errors and bounds for a particular aircraft is not proved here.

Nothing in this storage result provides arbitrary historical survivor answers after erasure. The task here is current-reference tracking and safe current control. Earlier retained distinctions may legitimately be discarded once raw operational eligibility expires, because this interface does not promise to reconstruct every later counterfactual query. The lower bounds in [1, 5] apply when that stronger query interface is demanded.

12 A parameter example and available checks

The supplied initial source reports an experimental package, but its checker, result manifest, metric input, and plotting CSV files are absent from the furnished directory. The original source and its reported counts remain in the review history. Those reports are not treated here as reproduced numerical evidence, and this source no longer depends on the missing files. The earlier supplied PDF remains unchanged in the archive. The matched v0.5 pair is the most recent compiled edition; the present source revision has not been compiled because PDF work is owner-paused.

12.1 An analytic design calculation

Consider six bodies assigned a 3×2 grid with nearest slot spacing 0.8 m. Choose

$$a = b = 1, \quad h = 1 \text{ s}, \quad H = 0.02 \text{ s}, \quad Q = 0.15 \text{ m}, \quad S = 0.10 \text{ m}, \quad D_b = 0.002 \text{ m/s}^2. \quad (55)$$

Take $B = Q + S = 0.25 \text{ m}$, $R = 0.252 \text{ m}$, body radius 0.04 m, $m = 2$, $D = 3$ and $\zeta = 0.4$. The radius obeys $R = B + D_b/(ab)$. Use profile rows in $\overline{B}_Q(0)$ and the ball projection above. This is a design calculation under the stipulated service and initial-state conditions, not a reported run.

For the critical plant,

$$\sigma_1 = e^{-1}(1 + \sqrt{2}) = 0.888139536 \dots, \quad \rho = \sigma_1 + 0.02 = 0.908139536 \dots < 1. \quad (56)$$

The conditional all-body bounds evaluate to center separation 0.296 m, speed 0.504 m/s, commanded acceleration 1.006 m/s², and effective acceleration 1.008 m/s². The separation exceeds the 0.08 m body diameter by 0.216 m. These values follow directly from theorem 4; they do not measure trajectories or establish that a particular communication system meets $H = 0.02 \text{ s}$.

12.2 A persistent-turnover workload with a smaller tracking certificate

The coarse bounds in (49)–(52) prove boundedness but need not improve on compactness of the cells. The following specified workload uses the same six-body physical parameters to give a smaller event-sensitive certificate. It is an analytical construction under the supplied services, not a simulation or a claim for arbitrary churn.

Take $\mathcal{S} = \overline{B}_{0.10}(0)$, exact fresh-reference delivery $s(i, k) = k$ and $\widehat{v}_{i,k} = \nu_k$, zero priority rounding, $m = 2$, $D = 3$, and now $\zeta = 1$. Give every physical role the two profile rows $-0.15u$ and $0.15u$ for a fixed unit vector u . Every current neutral row equals the supplied $\nu_k = (p_k, 1 - p_k)$, and every initial/startup priority row equals its then-reference. Thus every live $w_i(k) = \nu_k$, with $e_k = 0$, despite changes of the reference. Choose $p_k = 0.49$ initially and toggle it between 0.49 and 0.51 exactly at commits $k = 100n + 60$, $n \geq 0$. The reference jump has ℓ^1 size 0.04; the planned goal changes by 0.006 m only at those commits, and $\|q_{i,k}\| = 0.003$ m. Between them there is no neutral-reference change. Birth adds the current neutral row and close removes an equal row, so neither operation changes this mean.

Initially all six bodies are admitted and satisfy the cell initial-state premises. In cycle n , the FIFO service handles only body $i_n = 1 + (n \bmod 6)$. At the associated source stage σ_{100n+1} it selects physical exclusion and closes that body's old logical incarnation, using the ordered logical/source stages in section 2. The body completes its old physical command until commit $100n + 1$; physical exclusion and standby control begin at that commit. The standby goal $c^0 = r_{i_n} + q_{i_n}^- + s^-$ is chosen from the old held task and latch before close, bound in the source-stage decision, and held from the ensuing commit until readmission. Other bodies remain admitted. Use $P = 1$, $D_{\text{raw}} = 3$, $q_{\text{quiet}} = 5$, two alternating colours, and modulo size $M = 16$. After quiet time, the body receives a new incarnation and its current reference with bootstrap bound $B_{\text{boot}} = 1$. A chosen 50-second standby dwell and the remaining commit alignments fit inside the conservative 60-tick service envelope (45); the fresh readmission table is sampled for commit $100n + 60$. Its standby age is at least 58.98 seconds, hence exceeds 50 seconds. The core therefore has five or six bodies forever and never becomes fixed. At most one lifecycle request is outstanding, so this schedule respects the finite FIFO capacity. The direct-service production cap is $F = 84$ copies per epoch, with the same bounded frame, disposal and deadline premises as in section 11. The recurring role/reference schedule can use a 600-tick clock phase, a six-body cyclic selector and the current two-valued reference; the separate transport epoch uses modulo $M = 16$. The unbounded analysis index n need not be stored operationally.

Hold the latch at other commits. At readmission it may be held or recomputed if the settled-body guard permits, with an admissible projection residual at most 10^{-5} m. Use settled thresholds $R_g = V_g = 10^{-4}$ in their respective position/speed units. We do not require a forbidden recomputation when all admitted bodies are settled. Disturbances on every body may be any measurable signals of norm at most 0.002 m/s². Admit the single entrant only when its delayed certificate (38) is at most $a_* = 0.03$ m/s. The next calculation proves that this smaller gate passes on this particular schedule. It does not extend the broad-gate theorem to arbitrary small-threshold requests.

For the following standby calculation, τ denotes the numerical value of elapsed time in seconds; the displayed expressions give the numerical position and velocity bounds in metres and metres per second. For $\tau \geq 40$, the critical-plant homogeneous position and velocity contributions are bounded by

$$(0.502 + 1.006\tau)e^{-\tau}, \quad (0.502\tau + 0.504(\tau - 1))e^{-\tau}, \quad (57)$$

respectively. These bounds use $\|x(0) - c^0\| \leq R+B = 0.502$ and $\|v(0)\| \leq 0.504$. Variation of constants has position impulse kernel $\tau e^{-\tau}$, of integral one, and velocity kernel $(1 - \tau)e^{-\tau}$, of absolute integral $2/e$. The same norm bounds hold for arbitrary three-dimensional bounded forcing. Since $e^{-40} < 10^{-17}$ and both displayed envelopes decrease for $\tau \geq 40$, the source position error to c^0 is below 0.00201 m and source speed below 0.00148 m/s. Changing the planned goal at readmission adds at most 0.006 m.

The old latch has remained fixed throughout this standby interval, so

$$\mathfrak{a}_{i_n,k}^2 \leq (0.00801 + 0.504 \cdot 0.02)^2 + (0.00148 + 1.008 \cdot 0.02)^2 = 0.0007955377 < 0.03^2, \quad \mathfrak{w}_{i_n,k} < 0.022. \quad (58)$$

This closes the capture-time obligation for the declared schedule, including nonzero disturbance and observation age. The new planned reference is supplied before that source table and is held unchanged through commit.

Using σ_1 from (56), a conservative rational contraction bound is $\rho < \bar{\rho} = 0.909$. For every cell $U_k < 0.0049$ m/s, so the continuous-input part of J_k is at most $u = 1.02 \cdot 0.0049 = 0.004998$ m/s. At a readmission there are five continuing bodies. Their goal-change cost is at most $\sqrt{5} \cdot 0.006$ m/s, the admission cost is at most 0.03 m/s, the delayed entrant-speed contribution is below $0.02 \cdot 0.022$ m/s, and the projection residual contribution is at most $\sqrt{6} \cdot 10^{-5}$ m/s. Therefore the remaining event cost is below 0.04390 m/s. It is zero at other commits, including physical removal. Every 100-tick window consequently has $B_{\text{evt}} = 0.04390$ m/s. Applying corollary 1 with $\bar{\rho}$ gives

$$E_k \leq \bar{\rho}^k E_0 + \frac{0.004998}{1 - 0.909} + \frac{0.04390}{1 - 0.909^{100}} < \bar{\rho}^k E_0 + 0.099 \text{ m/s}. \quad (59)$$

Here the energy norm has speed units because $K = 1 \text{ s}^{-2}$. The certificate is independent of the completed-cycle count. In contrast, the cell bounds alone give

$$E_k \leq E_{\text{cell}} := \sqrt{6[K(R+B)^2 + V_{\text{cert}}^2]} = \sqrt{3.03612} \text{ m/s} > 1.74 \text{ m/s}. \quad (60)$$

Thus the eventual sampled error ceiling here is more than 17 times smaller than that cell-only ceiling. Between commits the eventual energy bound is below 0.104 m/s. Since $e_k = 0$ and $n_k \geq 5$, lemma 3 gives eventual pair-position task error below 0.148 m and common task error, relative to the actual latch, below 0.047 m. The corresponding cell-only pair bound is $2R = 0.504$ m for these shared profiles. Error relative to zero common translation still includes the separately allowed $S = 0.10$ m; it is not dropped from an absolute-position statement.

There is a nonempty parameter family behind this example. Keep the reserved geometry, gains, timing pattern and exact fresh-reference rule fixed, and replace the reference levels by $1/2 \pm \alpha$ with $0 < \alpha \leq 0.01$. The same certificate holds for $0 \leq H \leq 0.02$, $0 \leq D_b \leq 0.002$ and $0 \leq \delta_k \leq 10^{-5}$, retaining the same conservative constants. Positive delay, disturbance and projection error are allowed simultaneously. The strict gate and compactness-improvement margins above persist throughout this family. This supplies one quantitative persistent workload; it is not an optimization or evidence that every feasible workload has this accuracy.

12.3 Review diagnostics

The separately frozen review reports in the project's **review** directory identify newly executed bounded diagnostics and their exact scripts. ENTAIL checked 2,000 seeded critical-flow roster/projection events and 200 exact delayed-window steps. FOLIUM checked 4,000 critical-flow events and 8,820 finite-age/quarantine/modulo cases. These checks target formula instances and interface failures, use the project Python environment, and are not replays of the absent original experiment. Floating-point cases are not certified arithmetic or continuous-trajectory verification. The analytical arguments above retain responsibility for their universal conditional statements. A future end-to-end implementation must test actual source/live-state binding, table versions, workload/lifetime bounds, admission decisions, and every body's dynamics together.

13 Limitations, reproducibility obligations, and closing source review

This paper is a restricted infinite-execution base case under a trusted bounded coordination substrate and reserved workspace. Removing direct reference injection, weakening coherent tables or the commit clock, or replacing permanent cells are separate research questions. A successor must re-establish the affected guarantees when changing those contracts; they do not follow from this paper. The result is not a theorem for the original five programs run simultaneously. It supplies a new composition whose assumptions and modified updates are explicit. In particular, fixed task coordinates, reference injection, complete delayed physical tables, constrained latching and reserved cells cannot be removed while retaining the conclusion. The original support-conditioning problem, source-level simulator callbacks, and delayed relative-force controller keep their own source contracts.

The physically admitted set may change infinitely often, but the reference/commit grid limits how fast actual changes can occur. Several operations can share a commit, yet their costs are added under a fixed nonempty-core convention. The weighted event estimate is not a permission to execute a Zeno event sequence. The general lifecycle resource theorem is for accepted capacity-bounded requests with the broad gate, not arbitrary requests demanding arbitrarily small entry error. The smaller gate in section 12.2 has a separate proved capture-time margin for its stated workload. Greater geometric freedom, asynchronous uncoordinated physical commits, packet losses without an expiry/availability discipline, and heterogeneous gains require additional analysis.

The main error is relative to a declared moving task and an actual retained latch. Its vanishing is not historical restoration of every former preference, proof of agreement on a true proposition, or a grant of governance authority. Existing numerical influence can persist after raw messages expire. Quarantine closes a protocol path; it does not rewrite physical or informational history.

The prior-work comparisons concern open-system contraction, open formation, impulsive stability, and finite sequence reuse. The current reviews checked the primary open-system and formation sources and the protocol standards at their stated comparison scope. Initial retrieval of the impulsive-stability PDF failed; its alternate author-hosted route subsequently supplied the abstract and stated ISS/dwell-time framework, which support the limited contextual comparison. These are not proof dependencies or evidence that this combination is unprecedented.

Reproducibility boundary. The matched v0.5 source/PDF pair remains the latest compiled edition while PDF work is owner-paused. This v0.6 source-only revision has no matching new PDF yet. Earlier complete editions are archived, and the review history preserves the original source, separately frozen reports, and bounded checks. Missing original experiment/source-manifest files remain explicitly unavailable. The source is self-contained with respect to its TeX inputs. Compilation and page rendering are delivery checks and do not change the mathematical review’s scope. The five internal references are identified below by their actual revisions. A prior requested Markdown hash differs from the supplied revision; that unresolved historical discrepancy is recorded without treating either unseen bytes or a filename as a match. No manuscript submission is claimed.

AI-use disclosure and responsibility

This work was developed with assistance from AI agents operating within the HORTRAME Systems research architecture. Depending on the research stage, these agents contributed to literature analysis, mathematical formulation, proof reconstruction, adversarial review, counterexample search, simulation and code analysis, verification design, reproducibility checks, revision, and manuscript preparation. AI-agent outputs were treated as research artifacts subject to verification and correction rather than as

independent authority. G. Bolduc is the human author responsible for the manuscript, its claims, source selection, and publication decisions. AI-agent review does not constitute independent human peer review, proof-assistant certification, experimental validation, or publication acceptance.

The initial draft reports development by one AI assistant, including formulation, proof construction, source inspection, and numerical work; its original numerical companions are unavailable in this delivery. ENTAIL and FOLIUM subsequently produced separately frozen source reviews, compared findings, and revised assumptions, presentation, and evidence statements under the owner’s direction. They share mathematical sources and tooling, so their agreement is not independent human validation. No temporary subagents were used in this revision cycle. Guillaume Bolduc is identified as the human project author; final acceptance remains for his review. No proof-assistant certificate, physical experiment, or publication acceptance is claimed.

A Model changes and source binding

Source object	Retained mechanism	Changed or not inherited here
Historical restoration [1]	Delayed history matters; information and physical contracts remain separate.	Fixed task alphabet and injected reference replace the plain variable-support priority recurrence. No original endpoint is silently reused.
Disturbance recovery [2]	Equal-gain common/relative separation; exact local feedback.	Stale positions enter discrete latch recomputations, not the parent's continuously held centered-force decoder.
Sampled latch [3]	Fixed-cell energy contraction and nonexpansive position recentering.	Projection onto a convex translation set, physical roster changes, and delayed source tables are new. Exact source callback equivalence is not claimed.
Live admission shield [4]	Logical membership precedes physical influence; a fresh observation contract is necessary.	Pre-admission physical exclusion replaces the parent's free common-mode admission law; bounded cells and possible logical reference influence remain. No equivalence of shields is claimed. A robust source-stage gate is proved.
Finite history [5]	Retained input-dependent information is counted by the promised task.	Only current reference/control service is promised, not arbitrary historical conditional queries.

The supplied Markdown's actual SHA-256 is

65afd409ae6302ee6c20d1c435797d9de2f82389a4ef37f0f27256588ab51df2.

It does not match the earlier requested

37cb5798311a8d4c76a7adef7754c1f2f9ca9867438335dc9f9dbe2e93f16679.

This discrepancy is unresolved. The present proofs do not rely on an unseen version matching the requested hash.

B An exact held-input map for the nominal example

For $K = 1$, $C = 2$, held command c and force d , write $z = x - c - d$. For elapsed time $\tau \geq 0$,

$$x^+ = c + d + e^{-\tau}[(1 + \tau)z + \tau v], \quad (61)$$

$$v^+ = e^{-\tau}[-\tau z + (1 - \tau)v]. \quad (62)$$

Differentiation verifies the plant and the initial values. Splitting a cell into $h - H$ and H therefore gives its source and pre-commit states without interpolating between numerical integrator steps. For this critical plant and the energy with $K = 1$, the largest singular value at time h is

$$\sigma_h = e^{-h}(\sqrt{1 + h^2} + h). \quad (63)$$

This follows by calculating the two singular values of $e^{-h} \begin{bmatrix} 1+h & h \\ -h & 1-h \end{bmatrix}$. The formula makes the sufficient age limit $H < 1 - \sigma_h$ explicit in the normalized example. It is a sufficient condition for this proof, not an identified physical stability boundary.

C A dependency map for review

The raw-lifetime assumptions lead to theorem 1, not to erasure of influence. The reference service and bounded eligible ages lead to theorem 2, independently of a physical stability claim. The plant and coherent latch give lemmas 1 and 2 and then theorem 3. The separate all-body cell invariant proves theorem 4 before its bounds are used in theorem 5; this avoids assuming tracking safety to certify the admission that is supposed to establish tracking. The broad gate and a bounded accepted queue give theorem 6. Theorem 7 composes those specific results. The final convergence statement additionally requires the settled-body guard and a genuinely convergent, vanishing-input tail.

References

- [1] G. Bolduc, *Exact Restoration and Historical Memory in Delayed Priority Networks with Three-Dimensional Motion*, HORTRAME Research, supplied Markdown, v0.6, 3 Oct. 2026. Exact Markdown SHA-256 is stated in Appendix A.
- [2] G. Bolduc, *Translation-Equivariant Disturbance Recovery in Delayed Open Priority Networks*, supplied preprint, v05, 2026. Exact source SHA-256: 165c9b5dda01ae9573a2af28f0b9bd78ec380bce653e594cc4a3ba4662260773.
- [3] G. Bolduc, *A Sampled Common-Translation Latch for Nominal Double Integrators*, supplied addendum, v01, 6 Oct. 2026.
- [4] G. Bolduc, *A Common-Mode Shield for Live Admission in Delayed Open Priority Networks*, supplied addendum, v01, 6 Oct. 2026.
- [5] G. Bolduc, *Conditional Queries with Finite Historical States and Periodic Replacement*, supplied LaTeX manuscript, v09, 2026. Title and supplied revision are used without asserting publication.
- [6] M. Franceschelli and P. Frasca, “Stability of Open Multi-Agent Systems and Applications to Dynamic Consensus,” arXiv:1906.00890, 2019. Version-1 HTML, especially Sec. III and Theorem III.9, inspected. <https://arxiv.org/html/1906.00890v1>.
- [7] J. P. Hespanha, D. Liberzon, and A. R. Teel, “Lyapunov conditions for input-to-state stability of impulsive systems,” *Automatica*, vol. 44, no. 11, pp. 2735–2744, 2008. doi:10.1016/j.automatica.2008.03.021. Author-hosted abstract and ISS framework inspected for context. <https://web.ece.ucsb.edu/~hespanha/published/autom-impulsive-iss.pdf>.
- [8] P. Şekerciöglü, N. De Carli, and D. V. Dimarogonas, “Safe Formation Control of Open Multi-Robot Systems with Connectivity-Preserving Reconfiguration,” arXiv:2609.29939v1, 24 Sep. 2026. Secs. III–IV inspected. <https://arxiv.org/html/2609.29939v1>.
- [9] M. Gong, Y. Xia, and Y. Lv, “Time-varying formation for second-order open multi-agent systems with external disturbances,” *IEEE/CAA Journal of Automatica Sinica*, vol. 13, no. 6, pp. 1524–1526, June 2026. doi:10.1109/JAS.2025.125582. Bibliographic metadata retained from the initial draft; original-body and supplement comparison not independently established in this revision. Not a theorem dependency.
- [10] W. Eddy, ed., *Transmission Control Protocol (TCP)*, RFC 9293, Aug. 2022, Secs. 3.4.2–3.4.3. <https://www.rfc-editor.org/rfc/rfc9293.html>.
- [11] R. Elz and R. Bush, *Serial Number Arithmetic*, RFC 1982, Aug. 1996. <https://www.rfc-editor.org/rfc/rfc1982.html>.