

Exact Restoration and Historical Memory in Delayed Priority Networks with Three-Dimensional Motion

G. Bolduc
HORTRAME Research

Abstract—We study restoration of probability-valued priorities after finite changes in network membership. For fixed uniform or unequal directional communication delays, the endpoint depends on the complete initialized history. We derive the exact returned endpoint after one admission, finite mixing and proportional departure. Priority restoration is an affine equation in the admitted conditional preference; restoration of a three-dimensional objective centre is its projection through the anchor matrix. A causal packet protocol preserves origin-dependent historical transport, uses reception ordinals to order simultaneous readiness, and yields bounded-age stochastic dynamics after finite membership changes. Under bounded joint connectivity and positive current self-weights, those dynamics converge with an explicit sufficient contraction bound. Restoration over unknown admissible futures is characterized by finitely many endpoint functionals; a finite polynomial characterization additionally requires positive-polynomial initial and bootstrap data. We prove historical survivor-query memory bounds in distinct smooth-state and deterministic finite-bit models. A reference service retaining unmixed rows provides eventual restoration under a modified update law. Finally, a double-integrator controller converts priority convergence into point rendezvous or separated formation convergence. A staged finite-pool construction parks retired bodies in reserved finite-depth berths, while preserving all-body separation and speed and acceleration bounds. All conclusions retain their stated information, timing, storage and geometric assumptions.

Index Terms—networked control, consensus, communication delay, membership changes, historical information, formation control.

I. PROBLEM AND CONTRIBUTION

Consensus specifies agreement. Restoration specifies which agreed value must be recovered after an intervention. These properties differ when communication delays reweight historical rows or when departure conditions a probability row on its surviving support.

The delay-free restoration identity in the supplied Priority-State Dynamics manuscript [1, Theorem 4.1] motivates the present problem. Delayed stochastic consensus and delay-dependent endpoints are established subjects [2]–[4]. The fixed-delay invariant below is consistent with the compute/relay stationary weights in [3, Sec. IV-A]. The bounded-connectivity convergence theorem is a direct, explicit argument in a class narrower than the qualitative delay-robust consensus setting of [4]. Neither statement is presented as a first consensus result.

The article addresses exact support-map restoration, its three-dimensional projection, the information needed for historical survivor queries, and a causal realization connecting those

objects. Algebraic finite generation and the rank and code-counting methods used below are established mathematical tools. No first-publication claim is made for their application here. The supplied literature dossier remains a separate, attributable review; references whose relevant original statements have not been inspected are not used as theorem dependencies.

Related work also treats changing-dimensional optimization operators [5], storage for learning-unlearning tasks and ticketed information [6], collision/connectivity constraints in open double-integrator formations [7], and open-network coordination with dynamic links and processing delays [8]. These are comparisons with specified objectives and resources. The historical survivor-query interface and the staged physical construction below retain their own definitions; no predecessor theorem or first-publication claim is inferred from abstract-level comparison.

We distinguish three reference conventions. A *prescribed reference* is a fixed probability row supplied by the task. A *paired reference* is the endpoint of specified reference histories under a future admissible from both complete executions. A *historical survivor reference* is the mean of individually conditioned unmixed rows. No equality between these references is assumed.

A. Position relative to prior results

Tsianos and Rabbat [3, Sec. IV-A, eqs. (5)–(7)] derive stationary compute/relay weights in an augmented fixed-delay graph. Lemma 1 below specializes that mechanism to grouped history layers and unequal directional lags; neither augmentation nor a delay-dependent consensus value is claimed as new. Proskurnikov and Calafiore [4, Sec. IV] treat substantially weaker connectivity and weight regimes. Our stronger uniform assumptions buy an elementary common-column bound, not a more general consensus theorem.

Deplano et al. [5, Sec. II] already use labeled open operators. Our support map is probability conditioning, not a new theory of changing dimension. Hu et al. [8, Sec. 6.1] explicitly distinguish departing-soon and long-term remaining neighbors for their delayed quantized protocol. Our correct sequencer and bounded service assumptions are different resources and are not removed by calling a local update distributed.

Cherapanamjeri et al. [6, Secs. 1.1 and 4] distinguish central and ticketed deletion memory. The task below supplies survivor identities only, evaluates an average of rowwise conditioned original preferences, and counts all input-dependent retained information. No reduction from their learning task is asserted. The rank and covering arguments used here are classical

methods; the stated family, query map, and constants are the objects being proved.

Sekercioglu et al. [7, Secs. III–IV] use a formation manager, barrier-based dynamics, and switching conditions. Our invariant-cell and finite-parking construction has a different, more restrictive geometry. It is not a replacement for their reconfiguration theorem. The Stacks Project’s finite-generation result [9] is used only for an existential algebraic reduction, with its proof mechanism made explicit below. The initial and closing source-review record distinguishes full-text passages inspected in this revision from the supplied dossier’s abstract-only or unresolved entries. No first-publication or exhaustive-prior-art claim is made.

B. Result contracts and dependency structure

The paper contains distinct conditional results, not one algorithm with interchangeable states. The fixed-delay recurrence uses initialized delay registers. Theorem 1 uses an offline atomic transformation of every required history layer. The packet model instead transforms immutable received rows through a membership log and does not rewrite the physical past. Theorem 2 applies only after a proved fixed-support cut. Theorem 3 compares complete admissible futures common to the executions being compared. The historical query in Section V is a separate task defined on an unmixed table after an explicit erasure. Theorem 6 concerns an ideal double integrator, not a rigid-body flight engine.

The dependency chain is: stochasticity and age/connectivity hypotheses imply consensus; an invariant or future functional identifies its endpoint; a specified reference defines restoration; an anchor map transfers that endpoint to a goal; a separately proved controller transfers goal convergence to position. None of these arrows grants an unmodeled resource. In particular, Theorem 1 is not a correctness proof for the packet protocol, and neither controller below certifies the supplied flying-cat program.

All asymptotic claims assume finitely many support changes followed by an infinite fixed-support execution. All simulation statements in this revision must be bound to newly executable files or identified as inherited reports. Mathematical correctness and publication priority are separate questions.

II. PRIORITY STATES, DELAYS AND SUPPORT MAPS

For an ordered nonempty label set of size n , define

$$\begin{aligned}\Delta_n^+ &= \{p \in \mathbb{R}^{1 \times n} : p > 0, p\mathbf{1}_n = 1\}, \\ \mathcal{S}_n^+ &= \{W \in \mathbb{R}^{n \times n} : W > 0, W\mathbf{1}_n = \mathbf{1}_n\}.\end{aligned}\quad (1)$$

Here inequalities are entrywise. Unsubscripted vector norms are Euclidean; $\|\cdot\|_F$ is the Frobenius norm, and explicitly subscripted 1- and infinity-norms retain their usual meanings. A probability row has dimension $1 \times m$; stacking n such rows gives $W \in \mathbb{R}^{n \times m}$. The membership formulas specialize to $m = n$. Weight vectors $\kappa_s \in \mathbb{R}^n$ and $\pi_\sigma \in \mathbb{R}^M$ are columns; their transposes are the row functionals acting on stacked row tables. Only the full stacked κ and π_σ are normalized to total mass one. The anchor table has dimension $m \times 3$; physical vectors are identified with row coordinates when stacked in

anchor or position tables. We distinguish graph degree d_i , maximum lag D , physical cell diameter, and information dimension d by their local definitions.

Rows are preferences on objective labels. Physical dimension is independent of objective count. In the membership realization, each active incarnation contributes one objective; the fixed-support convergence algebra permits any finite objective count.

An arrival with incumbent insertion masses $\beta_i \in (0, 1)$, conditional newcomer preference $q \in \Delta_n^+$, and $\gamma \in (0, 1)$ has the table map

$$J(W) = \begin{bmatrix} \text{diag}(1 - \beta)W & \beta \\ \gamma q & 1 - \gamma \end{bmatrix}. \quad (2)$$

For a nonempty survivor set B , proportional departure maps an individual row to

$$D_B(p) = p_B/p(B), \quad p(B) = \sum_{j \in B} p_j > 0. \quad (3)$$

This is the unique probability normalization preserving all retained-coordinate ratios. It does not generally commute with mixing. Indeed, if rows have retained masses $s_1 \neq s_2$ and retained conditional rows $u \neq v$, then

$$\begin{aligned}D_B((p + q)/2) - (D_B(p) + D_B(q))/2 \\ = \frac{(s_1 - s_2)(u - v)}{2(s_1 + s_2)} \neq 0.\end{aligned}\quad (4)$$

Thus conditioning is not an affine kernel when $1 < |B| < n$. It is the identity for full support and a constant unit row for singleton support.

A. Fixed uniform and nonuniform delays

Let $n \geq 2$, and let G be a connected simple undirected graph with degrees d_i , $d_{\max} = \max_i d_i$, adjacency H , and degree matrix $D_G = \text{diag}(d_1, \dots, d_n)$. Both directions of each link exist. The directed delay from source j to receiver i is an integer $\tau_{ij} \in \{0, \dots, D\}$; opposite directions may have unequal delays. Uniform delay means $\tau_{ij} = d$ for all directed links. With $0 < c < 1/d_{\max}$, use

$$w_i(k + 1) = (1 - cd_i)w_i(k) + c \sum_{j: \{i, j\} \in E} w_j(k - \tau_{ij}). \quad (5)$$

Supply positive stochastic histories $W(-s)$, $0 \leq s \leq D$. Let H_s contain the adjacency entries of delay s , so $\sum_{s=0}^D H_s = H$. Put $B_0 = I - cD_G + cH_0$, $B_s = cH_s$ for $s > 0$, and let A be the companion matrix with first block row $(B_0, \dots, B_D)$ and identity shifts below. Then the actual recurrence state is

$$\begin{aligned}Y(k) &= \text{col}(W(k), \dots, W(k - D)), \\ Y(k + 1) &= AY(k).\end{aligned}\quad (6)$$

Lemma 1 (delayed endpoint). Every row remains positive stochastic, and $W(k) \rightarrow \mathbf{1}_n \mu$, where

$$\begin{aligned}
s_0 &= n + c \sum_{(j,i) \in E \rightarrow} \tau_{ij}, \\
\kappa_0^\top &= \mathbf{1}_n^\top / s_0, \\
\kappa_s^\top &= \frac{c}{s_0} \sum_{\ell=s}^D \mathbf{1}_n^\top H_\ell \quad (s \geq 1), \\
\mu &= \sum_{s=0}^D \kappa_s^\top W(-s) \in \Delta_n^+.
\end{aligned} \tag{7}$$

Here $E \rightarrow$ contains both directed orientations of each undirected edge. The κ entries are nonnegative and have total sum one. For $\eta = \min(c, \min_i(1 - cd_i))$, $L = n(D + 1)$, and $o_j = \max_{i, 0 \leq s \leq D} W_{ij}(-s) - \min_{i, 0 \leq s \leq D} W_{ij}(-s)$, a sufficient bound is

$$\begin{aligned}
\|W(k) - \mathbf{1}_n \mu\|_F &\leq \sqrt{n} \|o\|_2 \lambda^{\lfloor k/L \rfloor}, \\
\lambda &= 1 - \eta^L.
\end{aligned} \tag{8}$$

Proof. First, $B_0, \dots, B_D$ are nonnegative and their sum has unit row sums. Hence A is row stochastic, including its identity-shift rows. Every convex combination of strictly positive probability rows is strictly positive, even when some mixing coefficients vanish.

Second, fix one current-layer root. Read a row of A^L backward through the dependency graph. An initial history-layer row reaches its corresponding current layer in at most D shift steps. From a current node, one graph edge to a neighbor uses one read followed by at most D shifts. Connectivity supplies at most $n - 1$ such edges to the root. Thus at most $D + (n - 1)(D + 1) < L$ steps are needed. Pad the remaining steps at the root's positive current self-loop. Each nonzero read or self-loop is at least η , and shift coefficients are one. Every row therefore has at least η^L in that same root column. This argument does not require irreducibility of unused history registers.

Third, for a row-stochastic matrix F with $F_{ir} \geq \epsilon$ for all i , write

$$F = \epsilon \mathbf{1} e_r^\top + (1 - \epsilon) \tilde{F}, \tag{9}$$

where $\tilde{F}$ is row stochastic. For every scalar column x , its range satisfies

$$\max(Fx) - \min(Fx) \leq (1 - \epsilon)(\max x - \min x). \tag{10}$$

The common root term cancels between rows, while the residual values lie in the old scalar interval. Apply this to $F = A^L$, then use nonexpansion under intermediate stochastic steps. For each coordinate, nested minima and maxima have the same limit because their difference tends to zero. This proves consensus without an unjustified spectral-norm claim for a nonsymmetric companion matrix.

Finally, set $\kappa^\top = (\kappa_0^\top, \dots, \kappa_D^\top)$ and $\kappa_{D+1} = 0$. Direct block multiplication gives

$$\begin{aligned}
\kappa_0^\top B_0 + \kappa_1^\top &= \kappa_0^\top, \\
\kappa_0^\top B_s + \kappa_{s+1}^\top &= \kappa_s^\top \quad (1 \leq s \leq D).
\end{aligned} \tag{11}$$

The first identity uses $\mathbf{1}^\top H = \mathbf{1}^\top D_G$, not symmetry of the delay classes H_s . Each directed link with lag ℓ contributes to

exactly ℓ history layers, so $\sum_s \kappa_s^\top \mathbf{1} = 1$. Thus $\kappa^\top A = \kappa^\top$, and its conserved value identifies the consensus endpoint as (7). The endpoint lies in every nested coordinate interval; each of the n current rows is consequently within the remaining coordinate ranges. Summing their squared errors proves (8). Positive initial minima imply a positive limit. $\square$

With constant histories $W(-s) = W_0$, equation (7) becomes

$$\mu = \omega^\top W_0, \quad \omega_j = \frac{1 + c\psi_j}{n + c \sum_i \psi_i}, \quad \psi_j = \sum_{i: \{i,j\} \in E} \tau_{ij}. \tag{12}$$

The totals are *outgoing source delays*. For uniform delay d , $\psi_j = dd_j$. Uniform communication delay therefore need not give a uniform endpoint weight on an irregular graph. For unequal directional delays, uniform weights occur exactly when the outgoing totals are equal. A singleton holds its row, with $D = 0$, $A = [1]$, and no gain condition.

B. Boundary audit of the fixed-delay premise

For this Laplacian parametrization, $0 \leq c \leq 1/d_{\max}$ is exactly the nonnegativity range. Strict positivity of the diagonal requires $c < 1/d_{\max}$. Positive-row invariance itself only requires nonnegative unit-sum mixing rows. Consensus is a separate conclusion: $c = 0$ holds disagreeing rows; on K_2 with zero delay and $c = 1$, the matrix swaps them indefinitely. At zero delay on K_n , $c = 1/n$ gives the mean projector exactly and agreement after one update. Our common-column rate is sufficient, not an exact spectral rate. With delay, a larger real-valued stability range cannot be imported while retaining the probability interpretation.

Current-state equality is insufficient even with two nodes. Let $c = 1/4$, $\tau_{12} = 1$, $\tau_{21} = 0$, and $W(0)$ have both rows $(1/2, 1/2)$. Changing only row 2 of $W(-1)$ from $(3/4, 1/4)$ to $(1/4, 3/4)$ changes the endpoint from $(19/36, 17/36)$ to $(17/36, 19/36)$. The initial current disagreement is zero in both executions. Thus a current-row disagreement plot cannot replace the augmented-state endpoint calculation.

III. EXACT ROUND-TRIP RESTORATION AND ITS SPATIAL PROJECTION

The next theorem uses a prescribed atomic history map. It is distinct from the causal receiver protocol in Section 4. Initialized negative indices are delay registers, not claims that a newcomer physically existed in the past.

Supply incumbent histories $W_h \in \mathcal{S}_n^+$, $0 \leq h \leq D_*$, where $D_* = \max(D_f, D_e)$ covers final and enlarged delay depths. Join each history by (2), stack the resulting $(n + 1)$ -row snapshots into $\mathcal{V}_*$, and mix for a fixed stay ℓ . A nonnegative stochastic map K_s gives the enlarged snapshot used in final history layer s : $V(\ell - s) = K_s \mathcal{V}_*$. For a fixed enlarged recurrence, K_s is its companion power and current-layer selector when $\ell - s \geq 0$, or the initialized history selector otherwise. Initialize returned layer s by conditioning the incumbent rows of that snapshot on the old objectives.

Let $N = n + 1$ and $M_e = N(D_* + 1)$. Each K_s has size $N \times M_e$; only its first n rows are retained. Index initial stacked rows by (h, r) , and define

$$\begin{aligned}
a_{h,r} &= 1 - \beta_r \quad (r \leq n), \quad a_{h,N} = \gamma, \\
d_{s,i} &= \sum_{h,r} (K_s)_{i,(h,r)} a_{h,r}, \\
\theta_{h,r} &= a_{h,r} \sum_{s=0}^{D_f} \sum_{i=1}^n \kappa_{s,i} \frac{(K_s)_{i,(h,r)}}{d_{s,i}}, \quad \Theta = \sum_h \theta_{h,N}, \quad (13) \\
\mu_{\text{ref}} &= \sum_{s=0}^{D_f} \kappa_s^\top W_s, \quad g = \mu_{\text{ref}} - \sum_h \sum_{r=1}^n \theta_{h,r} (W_h)_{r,:}.
\end{aligned}$$

The reference and returned final epochs use the same graph, gain and delays. All histories, insertion masses, gains, maps, stay lengths and history policies are fixed independently of q .

This is an input-level oracle: the historical tables and the prescribed mixing maps are available to its designer. It does not assert that a node can choose q before learning an unknown future. Initializing or rewriting delay registers according to this contract is a counted intervention. The causal model below does not implement that intervention merely by receiving a departure record.

Theorem 1 (restoration with fixed delays). The returned endpoint satisfies

$$\mu_{\text{ret}} = \sum_h \sum_{r=1}^n \theta_{h,r} (W_h)_{r,:} + \Theta q, \quad \mu_{\text{ret}} - \mu_{\text{ref}} = \Theta q - g. \quad (14)$$

All $d_{s,i} \in (0, 1)$, $\theta_{h,r} \geq 0$, $\sum_{h,r} \theta_{h,r} = 1$, and $g\mathbf{1}_n = \Theta$. If $\Theta > 0$, a positive restoring row exists exactly when $g > 0$, and is uniquely $q_* = g/\Theta$. If $\Theta = 0$, restoration is equivalent to $g = 0$ and is independent of q .

For fixed anchors $z_j \in \mathbb{R}^3$, collected as rows of Z , the sampled local motion law

$$X(k+1) = (1 - \alpha)X(k) + \alpha W(k)Z, \quad 0 < \alpha < 2, \quad (15)$$

converges from every finite initial position to $\mathbf{1}_n \mu Z$. Surviving positions are retained at arrival and departure. Consequently the returned centre error is $(\Theta q - g)Z$. If $\Theta > 0$, positive spatial restoration is feasible exactly when

$$(g/\Theta)Z \in \text{relint conv}\{z_1, \dots, z_n\}. \quad (16)$$

When feasible, its solution is unique exactly when $\text{rank}[\mathbf{1}_n \ Z] = n$; otherwise it has infinitely many positive solutions. If $\Theta = 0$, spatial restoration is equivalent to $gZ = 0$.

Proof. The old-objective part of initial row (h, r) is $a_{h,r} C_{h,r}$, where $C_{h,r} = (W_h)_{r,:}$ for incumbents and $C_{h,N} = q$. Every $a_{h,r} \in (0, 1)$; stochastic mixing therefore gives $d_{s,i} \in (0, 1)$. The conditioned returned row is

$$(U_s)_{i,:} = \frac{\sum_{h,r} (K_s)_{i,(h,r)} a_{h,r} C_{h,r}}{d_{s,i}}. \quad (17)$$

It is positive stochastic. Substitute (17) into (7) and interchange finite sums. This gives $\mu_{\text{ret}} = \sum_{h,r} \theta_{h,r} C_{h,r}$. Summing coefficients cancels their denominators and gives one, proving (14) and $g\mathbf{1}_n = \Theta$. Solving $\Theta q = g$ gives both priority cases.

Set $E(k) = X(k) - \mathbf{1}_n \mu Z$, $Q(k) = W(k) - \mathbf{1}_n \mu$. Iterating (15) gives

$$E(k) = (1 - \alpha)^k E(0) + \alpha \sum_{t=0}^{k-1} (1 - \alpha)^{k-1-t} Q(t)Z. \quad (18)$$

The first term vanishes. Since $|1 - \alpha| < 1$ and $Q(t)Z \rightarrow 0$, split the sum at a fixed large index: its finite prefix decays, and its tail is bounded by an arbitrarily small input bound times the convergent geometric sum. Thus $E(k) \rightarrow 0$. A finite stay with a finite anchor table and finite gains yields finite positions by finite induction, so this proof applies to positions actually retained at departure. It makes no collision or actuator assertion.

The image of Δ_n^+ under $q \mapsto qZ$ is the relative interior of the anchor hull. For one inclusion, the zero-sum row space maps onto the span of anchor differences; a finite-dimensional right inverse sends a sufficiently small relative neighborhood to positive rows. For the reverse inclusion, let $\bar{z} = n^{-1} \sum_j z_j$. For any relative-interior y , sufficiently small $e > 0$ makes $y' = (y - e\bar{z})/(1 - e)$ belong to the hull. Write $y' = pZ$, $p \geq 0$, $p\mathbf{1} = 1$; then $(1 - e)p + (e/n)\mathbf{1}^\top$ is positive and maps to y . If $y = \bar{z}$, use the uniform row directly. This proves (16).

Two restoring probability rows differ by a row h with $h\mathbf{1} = 0$, $hZ = 0$. Such a nonzero row exists exactly when the displayed rank is less than n . Perturbing a positive solution by sufficiently small th yields infinitely many solutions. Full rank gives uniqueness. Since the augmented matrix has four columns, uniqueness requires $n \leq 4$. $\square$

At zero delay there is one snapshot and $\kappa_{0,i} = 1/n$. Formula (13) reduces to $d_i = \sum_r T_{ir} a_r$, $\theta_r = (a_r/n) \sum_i T_{ir}/d_i$, and $q_* = (n^{-1} \mathbf{1}^\top W - \sum_{r \leq n} \theta_r w_r)/\theta_{n+1}$ when the denominator is positive: the supplied Theorem 4.1 [1]. Equation (14) extends its priority identity; equations (15)–(16) give its three-dimensional consequence.

A. Conditioning and implementation sensitivity

Positive retained mass guarantees that a quotient exists, not that it is uniformly well conditioned. For positive rows u, v and nonempty B ,

$$\|D_B(u) - D_B(v)\|_1 \leq \frac{2\|u - v\|_1}{u(B)}. \quad (19)$$

Indeed, add and subtract $v_B/u(B)$; the two terms are bounded by $\|u_B - v_B\|_1/u(B)$ and $|u(B) - v(B)|/u(B)$. For $u = (3\delta/4, \delta/4, 1 - \delta)$ and $v = (\delta/4, 3\delta/4, 1 - \delta)$, the input distance is δ while departure to $B = \{1, 2\}$ produces distance one. No mass-independent numerical robustness follows.

The restoration calculation has a second sensitivity. Suppose $\Theta > 0$, $q_* = g/\Theta$, $\hat{g} = g + e_g$ and $\hat{\Theta} = \Theta + e_\Theta$, with $|e_\Theta| < \Theta$. Direct subtraction gives

$$\|\hat{q} - q_*\|_1 \leq \frac{\|e_g\|_1 + |e_\Theta| \|q_*\|_1}{\Theta - |e_\Theta|}. \quad (20)$$

For a feasible q_* , its norm is one. Positivity of the computed row is guaranteed when its coordinate error is below $\min_j q_{*,j}$. Normalization must also be checked: arbitrary arithmetic errors

need not satisfy $\hat{g}\mathbf{1} = \hat{\Theta}$. This parameter-selection sensitivity is different from the forward endpoint error $\Theta(q - q_*)$. Clipping negative entries and renormalizing defines a different candidate row whose residual must be recomputed. The exact theorem itself uses no floating-point repair.

IV. CAUSAL HISTORICAL PACKETS AND VARYING DELAYS

A. Protocol and service assumptions

Actions have a locally finite total order indexed by k . State $w_i(k)$ means the stored row immediately before primitive action k ; its result is $w_i(k+1)$. A send at g snapshots $w_j(g)$. A later consume at k reads that immutable snapshot, so its lag is $k - g$. Atomic batches are defined as single primitives only under a processing bound for the entire batch. An ordered expansion into subactions is also permitted, but all constants must then be recomputed. This is an analysis interleaving, not a clock available to nodes. A specified correct sequencer emits a finite ordered membership log. Every record supplies its predecessor, ordered support, objective metadata and row map. Arrival masses are a declared total function of immutable row origin, including a positive default for retired origins. A stored packet retains its original sender's origin through all subsequent maps. Departure uses (3). Maps along the log compose in order; endpoint supports alone do not determine their composition. For example, starting with $(1/4, 3/4)$, adding a label with mass $1/2$ and then removing the second old label gives $(1/5, 4/5)$. Reversing those operations gives $(1/2, 1/2)$ on the same final labels.

Each destination incarnation stores its applied prefix, current positive row when initialized, pending records, latest source buffer, largest accepted source sequence and a ready-origin FIFO. A newborn waits without a priority row, readiness or sends until a coherent positive bootstrap is applied. It first applies known contiguous successors. Applying its own removal record retires it locally, clears its buffers/FIFO and stops priority operations. Earlier packets from that source remain eligible at continuing destinations; later packets addressed to a retired or different destination incarnation are discarded.

A valid new data copy receives a unique increasing *receiver-local reception ordinal*. Its immutable payload names a source row, coherent source prefix/support, source and destination incarnations, sequence and generation identity. Invalid rows, wrong destinations, own-origin copies and duplicate or nonnew sequences are rejected. Conflicting payloads with one identity are excluded by the compliant model. Future-prefix copies wait. Each destination incarnation initializes its accepted-reception counter to zero. Refresh readiness after every accepted data reception, completed coherent bootstrap and completed contiguous-record application batch, and before every priority activation. Before activation, apply available contiguous records and refresh; if the receiver has retired, perform no priority update. No readiness exists before coherent bootstrap. At each readiness refresh, append newly eligible retained buffers in increasing ordinal of their current copies; already-ready origins retain position under eligible replacement. A future-prefix replacement removes its origin until it becomes eligible again. A contiguous record batch finishes before refresh unless

it retires the receiver. On activation, consume the first ready buffer, transport its immutable row through the known log, and mix it with the current row. If none is ready, hold. A nontrivial consume uses a declared weight in $(0, 1)$; quantitative convergence additionally requires the uniform lower bounds specified below. A send snapshots the actual current row to a destination in its applied support. Simultaneous physical-time primitives require a declared internal order. Records and nodes are assumed correct; this protocol does not solve Byzantine membership agreement.

Assume each relevant record and retirement is completely applied within P actions of its emission. A newborn's coherent bootstrap must be available and completely applied within the same bound measured from emission of its birth record, not from a separately delayed bootstrap emission. A newborn retired before initialization instead completes terminal retirement within that birth-bound deadline and disposes of its retained copies; data to continuing receivers is delivered within D_{net} actions of generation; and a ready retained copy is consumed or explicitly superseded within Q more actions. These are correct-channel, processing and load contracts, with positive integer bounds. A deadline $e + T$ refers to the state before action $e + T$, when the bounded operation has already taken effect. With at most J retained ready origins and one receiver activation in each A_s -action window, $Q = A_s J$ is a sufficient service bound. Here J counts still-buffered retired origins as well as current ones. Eligible replacement keeps a ready origin in place; a replacement that is not yet eligible is a different readiness interval. Bounded delivery without bounded readiness/service is insufficient.

Lemma 2 (causal fixed-support cut). Every current active row and consumed transported copy is positive normalized on its local applied support. Every consumed generation has age at most $D = \max(P, D_{\text{net}}) + Q$. If the last record is emitted at k_R , then $K_0 = k_R + P + 2D + 2$ is a sufficient analysis cut at which the needed D -layer histories and all subsequently used generations have final support.

Proof. Arrival, conditioning and convex mixing preserve positivity and normalization. At a use, all required maps are committed, received and known; no future state or event is used. At generation g , every record needed for the source prefix, including the destination birth, has already been emitted. Both that prefix and the destination bootstrap are therefore applied by $g + P$. No additional unbounded bootstrap-wait clock remains. Delivery takes D_{net} , then service takes Q , giving the age bound. All final record applications finish by $k_R + P$. Generations before that bound cannot be consumed after $k_R + P + D + 1$. At K_0 , even the oldest required history precedes the cut by only D and lies after that index. The two extra actions avoid pre/post-action ambiguity. No node detects this cut or resets its state there. $\square$

A separate bootstrap-emission deadline is insufficient. For example, let birth occur at action 0, a packet be generated at 1 and received at 2, and a separately emitted bootstrap at 100 be applied at 101 before consumption at 102. Bounds $P = 10$, $D_{\text{net}} = 1$ and $Q = 1$ hold relative to that bootstrap's own emission, but the consumed packet has age 101, not at most 11. Binding bootstrap completion to the birth record excludes

this execution. This is an assumption boundary, not a property obtained from FIFO alone.

B. Contraction and restoration over admissible futures

After the cut, write the actual recurrence as

$$w_i(k+1) = a_{ii}(k)w_i(k) + \sum_{j \neq i} a_{ij}(k)w_j(k - \tau_{ij}(k)). \quad (21)$$

Each row of coefficients is stochastic; $a_{ii} \geq \eta > 0$; every positive off-diagonal coefficient is at least η , with $\eta < 1$; and $0 \leq \tau_{ij} \leq D$. An inactive row holds. The union of actual used influence edges in every $B \geq 1$ consecutive indices is strongly connected. Sends and receptions that hold rows are included in the index. Uniform, unequal directional and bounded time-varying delays are allowed.

Theorem 2 (asynchronous endpoint). For $n \geq 2$, let $M = n(D+1)$, $L = (n-1)(B+D) + D$, and $\lambda = 1 - \eta^L$. Every length- L augmented product has every entry of any chosen initial current-root column at least η^L . Consequently

$$\|W(k) - \mathbf{1}_n \mu_\sigma\|_F \leq \sqrt{n} \|o\|_2 \lambda^{\lfloor (k-K_0)/L \rfloor}, \quad (22)$$

$$\mu_\sigma = \pi_\sigma^\top Y(K_0), \quad k \geq K_0.$$

Here $o_j = \max_{i,s} w_{ij}(K_0 - s) - \min_{i,s} w_{ij}(K_0 - s)$, with $0 \leq s \leq D$, and π_σ is a nonnegative unit-sum column determined by the complete admissible future σ . Augmented row hulls are nested compact outer endpoint sets with intersection $\{\mu_\sigma\}$. They need not equal sets of reachable endpoints.

Proof. Track one initial current-root coefficient. Own-state weight preserves its positive lower bound at least η^{v-t} at later time v . For an informed set present at time u , wait D steps. Every influence from that set in the next B steps then reads an informed time. Strong connectivity of their union provides an edge to the complement unless all nodes are informed. Its weight transmits at least η^{h+1-t} to the receiver by time $h+1$. Thus one more node becomes informed within $D+B$ steps. Repeat at most $n-1$ times; an additional D shifts inform every history layer. This proves the common-column bound.

Subtract that column as in Lemma 1 to contract augmented coordinate ranges by λ . Minima and maxima are nested; positive initial minima persist. This proves the positive common endpoint and (22). Apply the same scalar argument to each column of the product initialized with the identity. All its rows converge to one nonnegative unit-sum row $\pi_\sigma^\top$, proving the endpoint formula. Each new augmented row is a convex combination of existing rows; their hulls are nested, contain every later row and its limit, and have vanishing coordinate diameters. $\square$

For a singleton the current row holds and histories fill in at most D shifts; do not apply a formula with $L = 0$. The contraction concerns augmented ranges. Current-row ranges alone need not be monotone.

Theorem 3 (unknown futures). Fix a nonempty family of actual futures independent of the numerical parameter q , including send, receive, buffer, coefficient and lag choices. Let $\Pi = \{\pi_\sigma\}$, choose any $\pi_0 \in \Pi$, and let $Y(q)$ be its terminal

actual history. For a prescribed row ν , restoration for every future is equivalent to

$$\begin{aligned} \pi_0^\top Y(q) &= \nu, \\ h^\top Y(q) &= 0 \quad (h \in \mathcal{H}), \\ \mathcal{H} &= \text{span}\{\pi - \pi_0 : \pi \in \Pi\}. \end{aligned} \quad (23)$$

At most M actual futures determine these equations. For paired reference histories, the exact condition is $\pi^\top(Y(q) - Y_{\text{ref}}) = 0$ for all $\pi \in \Pi$, provided those futures are admissible from both complete executions, including pending packet availability. A basis drawn from at most M such futures suffices. Spatial restoration multiplies these expressions by Z .

Proof. Subtract the base endpoint equation. The difference span has dimension at most $M-1$, since all π sum to one; select a basis from its actual spanning differences. Include the base future. The paired formula follows from a common endpoint functional, then a basis of $\text{span } \Pi$. Abstract lag choices cannot supply nonexistent initial packets, which is why common admissibility is required. $\square$

This is an existential finite characterization, not an enumeration algorithm for all futures. Computing π_σ from an arbitrary infinite execution is not supplied by the span argument. With supplied finite-period rational recurrences, endpoint functionals can instead be computed from their period matrices, provided their consensus hypotheses are established. A schedule-independent mean is not guaranteed: alternating zero-lag single-row half-weight updates give invariant rows $(1/3, 2/3)$ or $(2/3, 1/3)$ when their order is reversed.

C. Finite polynomial certificates and their scope

Suppose every parameter-dependent initial/bootstrap row has coordinates $P_j(q)/\sum_h P_h(q)$, with nonnegative polynomial coefficients and positive numerators and denominators on the positive parameter simplex. All event masses, weights and meta-data policies are fixed independently of q . Conditioning restricts its numerator list; arrival changes it to $((1-\beta)P, \beta \sum P)$; and a fixed convex mixture of P/S and Q/T has numerator $(1-\alpha)P_j T + \alpha Q_j S$ and denominator ST . Induction therefore gives positive rational terminal histories. Clearing their positive denominators makes (23) and the paired or spatial conditions polynomial equations.

For any family of finite departure prefixes in a fixed number of parameter variables, finitely many original equations determine their common restoration set. To see this without a stopping claim, every infinite sequence of nonnegative integer exponent vectors contains a componentwise nondecreasing pair: extract an infinite constant or increasing subsequence in each coordinate in turn. Thus an upward-closed exponent set has finitely many minimal elements. In a graded lexicographic monomial order, the leading exponents of an ideal form such a set. Choose polynomials realizing its minimal elements. Repeated leading-term cancellation by their monomial multiples terminates by well-ordering and generates the ideal. Each chosen generator is a finite combination of original restoration equations; the union of their prefix indices is a finite determining family. This is the Hilbert basis argument [9]. It gives no uniform size bound or general stopping algorithm.

For the fixed-delay snapshot round trip, an explicit finite bound is also available. Pad its enlarged companion matrix A_e to depth D_* , of size $M_e = N(D_* + 1)$. Let a stack the old-objective masses in (13) and let $v_j(q)$ stack incumbent coordinate masses $(1 - \beta_r)(W_h)_{r,j}$ and newcomer masses γq_j . Enumerate the p positive final weights as (r_t, λ_t) , where r_t selects the relevant incumbent history row. Set

$$T = A_e^{\otimes p}, \quad K = M_e^p, \quad \rho^\top = e_{r_1}^\top \otimes \cdots \otimes e_{r_p}^\top, \\ \zeta_j(q) = \sum_{t=1}^p \lambda_t a^{\otimes(t-1)} \otimes v_j(q) \otimes a^{\otimes(p-t)} - \mu_{\text{ref},j} a^{\otimes p}. \quad (24)$$

Restoration for every stay $\ell \geq 0$ is equivalent to $\rho^\top T^\ell \zeta_j(q) = 0$ for $0 \leq \ell < K$, $j < n$, together with the positive simplex constraints. Indeed multiplying the endpoint residual by $\prod_t e_{r_t}^\top A_e^\ell a > 0$ gives this scalar tensor expression. Cayley–Hamilton supplies an order- K recurrence, so vanishing of its first K values implies vanishing of all values. The equations are affine in q , despite the potentially large bound. Explicit storage of T would require M_e^{2p} entries; the existence of the certificate is not a tractability theorem. One may use a smaller proved annihilating polynomial or a reachable Krylov subspace in a concrete instance, but no uniform efficient procedure is claimed here. Projecting the ζ_j through Z gives the corresponding spatial certificate.

The positive-polynomial premise cannot be dropped. For example, allow a positive bootstrap with old conditional first coordinate $f(p) = 1/2 + \sin(1/p)/4$, $0 < p < 1$, and self-mass $1/2$. Take two uniform incumbent rows, give both insertion mass $1/2$, and bootstrap the newcomer as $(f/2, (1-f)/2, 1/2)$. Remove old label 2. The remaining conditional rows are $(1/3, 2/3)$ and $(f/(f+1), 1/(f+1))$. Then use fresh half-weight exchanges in the order with invariant weights $(2/3, 1/3)$. Their endpoint first-coordinate error from target $1/3$ is $(2f-1)/(9(f+1))$. Restoration occurs exactly at $p = 1/(k\pi)$, $k \geq 1$: an infinite proper discrete set that no finite polynomial system on this one-dimensional simplex describes. Linear endpoint functionals remain valid for this bootstrap; the polynomial conclusion does not.

D. An actual timing obstruction

Use incumbent rows $u_a = (3/4, 1/4)$, $u_b = (1/2, 1/2)$. Admit c with insertion masses $1/3$ at origin a , $2/3$ at origin b , and bootstrap $(2p/3, 2(1-p)/3, 1/3)$. Before a learns arrival, it consumes one initial b packet with weight $1/2$, obtaining $(5/8, 3/8)$. It buffers a future-prefix packet from c , then another prefix-zero b packet. Applying arrival and consuming the already-ready b packet gives $(7/24, 5/24, 1/2)$, because that packet transports using origin b 's mass. The c copy is ready behind it.

Remove c . Branch E applies removal before consuming its copy with weight $3/4$; branch L consumes first and applies removal second. The returned first coordinates at a are

$$s_E(p) = 7/48 + 3p/4, \quad s_L(p) = 7/60 + 4p/5. \quad (25)$$

Node b stores removal until arrival is received, then applies both and returns to $(1/2, 1/2)$. Node c retires. The finite prefix has bounded record, delivery and FIFO service times. Continue with actual fresh sends, receptions and half-weight activations, in order F ($a \rightarrow b$, then updated $b \rightarrow a$) or R (the reverse). Their period matrices and invariant rows are

$$C_F = \begin{bmatrix} 3/4 & 1/4 \\ 1/2 & 1/2 \end{bmatrix}, \quad \pi_F^\top = (2/3, 1/3), \\ C_R = \begin{bmatrix} 1/2 & 1/2 \\ 1/4 & 3/4 \end{bmatrix}, \quad \pi_R^\top = (1/3, 2/3). \quad (26)$$

For the prescribed initial mean $(5/8, 3/8)$, branch E needs $p = 13/18$ under F and $p = 35/36$ under R. Even if the entire positive bootstrap is selectable, F requires returned a -coordinate $11/16$ and R requires $7/8$, so no one bootstrap restores that prescribed target under both futures.

For the *different* reference obtained by running the original incumbent table under F, the endpoint first coordinate is $2/3$. Restoration requires $s_E = 3/4$ or $s_L = 3/4$, hence respectively $p = 29/36$ and $p = 19/24$. This departure-placement obstruction concerns fixed self-mass $1/3$. It does not extend to freely selectable self-mass: bootstrap $(29/72, 7/72, 1/2)$ restores both branches. With distinct anchors, each residual row $(h, -h)$ produces a nonzero centre error $h(z_a - z_b)$. These examples use generated packets, not fictitious lag entries.

V. HISTORICAL SURVIVOR-QUERY MEMORY

Capture an unmixed positive table W^0 . A query $\emptyset \neq B \subseteq \{1, \dots, n\}$ asks for

$$\nu_B(W^0) = \frac{1}{|B|} \sum_{i \in B} D_B(W_i^0). \quad (27)$$

At the erasing observation cut all current rows are uniform, ordinary pending packets/buffers are empty, and nonhistorical labels, clocks and scheduling data are input-independent. Every remaining input-dependent register, physical state, external history or service counts as historical memory. The query supplies survivor labels, not deleted numerical data. The smooth-state and finite-bit models below are distinct.

A. Erasure and smooth-state necessity

Complete averaging can erase the query answer. The positive table

$$W^0 = \begin{bmatrix} 1/2 & 1/4 & 1/4 \\ 1/4 & 1/4 & 1/2 \\ 1/4 & 1/2 & 1/4 \end{bmatrix} \quad (28)$$

and the uniform table have the same full mean. After complete averaging their current tables coincide, whereas their $B = \{1, 2\}$ answers are $(7/12, 5/12)$ and $(1/2, 1/2)$. The averaging is realizable: send every initial row before any update; at each receiver consume the other initial snapshots with weights $1/2, 1/3, \dots, 1/n$. Induction gives the mean of the snapshots consumed so far. After all consumptions and finite drain, ordinary state is uniform. Any retained originals must be counted separately.

Theorem 4 (smooth historical memory). Let $n \geq 3$, $d = (n-1)(n-2)/2$, $h = 1/\lceil 4n(n-2) \rceil$. There is a positive doubly stochastic family $W(x)$, $x \in (-h, h)^d$, with identical complete averages, for which exact answers to all queries (27) require at least d historical real coordinates under a C^1 encoder and C^1 query decoders on neighborhoods of the encoder image. Storing d coordinates attains this bound for the family; storing $n(n-1)$ independent row coordinates answers arbitrary positive tables.

Proof. Set $W(x) = \mathbf{1}\mathbf{1}^\top/n + \Delta(x)$, where Δ is skew-symmetric, has free entries $\Delta_{ij} = x_{ij}$ for $i < j < n$, and $\Delta_{in} = -\sum_{j < n} \Delta_{ij}$. All row sums, hence column sums, vanish. Last-column magnitudes are at most $(n-2)h = 1/(4n)$; every entry of W is at least $3/(4n)$. For query $B = \{i, j\}$, $i < j < n$, its first coordinate is

$$f(nx_{ij}) = \frac{1}{2} - \frac{nx_{ij}}{4 - (nx_{ij})^2}, \quad f'(t) = -\frac{4+t^2}{(4-t^2)^2} < 0. \quad (29)$$

The d free-pair answers have a diagonal derivative of rank d . If they factor through an encoder into r real coordinates, the chain rule bounds that rank by r . Thus $r \geq d$. The free parameters reconstruct the table and all its queries; full row storage gives the stated general upper representation. $\square$

The theorem excludes discontinuous digit encodings into one exact real. It concerns historical survivor queries after the declared erasure, not one known endpoint or a global network-storage minimum. A known invertible finite mixing matrix can permit exact reconstruction; finite-precision inversion may be unstable. For $P = (1-\alpha)I + \alpha\mathbf{1}\mathbf{1}^\top/n$, $0 < \alpha < 1$, its inverse on disagreement amplifies perturbations by $(1-\alpha)^{-k}$ at time k .

B. Deterministic bit-error bounds

Use the closed parameter box $[-h, h]^d$, and define

$$t_0 = nh = 1/\lceil 4(n-2) \rceil, \quad \ell_0 = f(-t_0) - f(t_0) = \frac{2t_0}{4 - t_0^2}. \quad (30)$$

Define the loss precisely as

$$E = \sup_{x \in [-h, h]^d} \max_{\emptyset \neq B \subseteq [n]} \|\hat{\nu}_B(\text{enc}(x)) - \nu_B(W(x))\|_2. \quad (31)$$

The maximum is over individual queries, not an Euclidean norm of the concatenation of every possible answer. The bit budget b is a nonnegative integer, and each decoder is fixed before the input is chosen.

Theorem 5 (historical bits). Any deterministic encoder into at most 2^b historical states, with arbitrary query decoders, has worst Euclidean query error

$$E \geq \frac{\ell_0}{\sqrt{2}} 2^{-b/d}. \quad (32)$$

A product midpoint code attains this bound on the free-pair queries when $2^b = K^d$. On all survivor queries it attains the same exponent, with upper bound $(32(n-1)(n-2)/3)\ell_0 2^{-b/d}$. For $n = 3$, the exact minimax error over all queries is $8 \cdot 2^{-b}/(63\sqrt{2})$.

Proof. Free-pair first coordinates range independently through a box of side ℓ_0 . Orthogonally projecting each decoded vector onto the probability line $(y, 1-y)$ cannot increase its error. Pair error at most E implies first-coordinate error at most $E/\sqrt{2}$. Each decoder state therefore covers at most volume $(\sqrt{2}E)^d$ in the query-coordinate box. The 2^b possible states cover its volume ℓ_0^d , proving (32), without an encoder regularity premise.

Divide each query-coordinate interval into K equal cells and encode their d indices. Midpoint pair error is at most $\ell_0/(\sqrt{2}K)$. Inverting (29) reconstructs each free parameter to error at most $2\ell_0/(nK)$, since $|f'| \geq 1/4$. The reconstructed table retains the same positivity bound. Its row ℓ^1 errors are at most $(n-1)(n-2)2\ell_0/(nK)$. For positive rows,

$$\|D_B(u) - D_B(v)\|_1 \leq \frac{2\|u - v\|_1}{\sum_{j \in B} u_j}. \quad (33)$$

The denominator is at least $3/(4n)$, so every averaged query has Euclidean error at most $16(n-1)(n-2)\ell_0/(3K)$. Taking $K = \lceil 2^{b/d} \rceil \geq 2^{b/d}/2$ proves the upper bound. For $n = 3$, all three pair coordinates are $y = f(3x)$ or $1-y$, the full-set mean is fixed and singleton answers are exact. Quantizing this one interval of length $8/63$ into 2^b cells attains (32) for every query. $\square$

Only retained input-dependent information is counted as historical bits. Public constants and decoder code are fixed independently of the initial table. Computation, arithmetic representation and log/identifier storage remain separate resource costs.

C. Reference-assisted eventual restoration

A specified correct stable service captures each neutral row before erasure. It applies the declared origin-specific arrival maps and survivor conditioning to those rows, inserts newborn neutral rows and removes retired rows; communication mixing never changes the cache. It computes the current neutral-row mean ν_r , using $O(n_r^2)$ real row values plus separately counted control/log metadata. It delivers prefix-labelled summaries with bounded receipt/application. Future summaries wait. Until a compatible summary exists, the receiver uses a specified positive placeholder without claiming a reference-error bound.

After finitely many records, bounded final-summary application and the actual-history cut, let ν be the final summary. Modify each active update to

$$\begin{aligned} w_i(k+1) &= b_i w_i(k) + \sum_{j \neq i} a_{ij} w_j(k - \tau_{ij}) + \zeta_i \nu, \\ b_i + \sum_{j \neq i} a_{ij} + \zeta_i &= 1. \end{aligned} \quad (34)$$

with nonnegative coefficients, $\zeta_i \geq \zeta_0 \in (0, 1)$, used ages at most D , and each node active once per A_u actions. Then, from a sufficiently late analysis index K ,

$$\begin{aligned} E(k) &\leq E(K)(1 - \zeta_0)^{\lfloor (k-K)/(A_u+D) \rfloor}, \\ E(k) &= \max_{i, 0 \leq s \leq D} \|w_i(k-s) - \nu\|_\infty. \end{aligned} \quad (35)$$

To prove this, subtract v . Every active update has error at most $(1 - \zeta_0)E(k)$; holds and shifts cannot increase the augmented maximum. Within A_u actions each current row has this improved bound relative to a block's starting maximum. Later updates retain that same block bound. After D more shifts all history rows have it, proving (35) by blocks.

With no exogenous change to an original neutral preference, all original incumbents surviving, all newcomers eventually leaving, and only finitely many changes occurring, each original neutral row returns to its initial row: insertion scales its original coordinates together and later conditioning preserves their ratios. Hence the final summary is the original incumbent mean. This is eventual restoration under the new service and update law. The service is a counted information resource; (34) is not a conservation property of the plain protocol or an implementation of fault-tolerant membership agreement.

VI. CONTINUOUS THREE-DIMENSIONAL MOTION

A. Convex-cell filter and endpoint lift

Use exact local feedback in $\dot{x}_i = v_i$, $\dot{v}_i = u_i$, with $x_i \in \mathbb{R}^3$. For a nonempty compact convex cell C_i of diameter d_i , bounded measurable goal $c_i(t) \in C_i$, and gains $a_i, b_i > 0$, set

$$u_i = a_i b_i (c_i - x_i) - (a_i + b_i) v_i, \quad y_i = x_i + v_i / b_i. \quad (36)$$

Theorem 6 (continuous lift). If $x_i(0), y_i(0) \in C_i$, then $x_i, y_i \in C_i$ for all time, $\|v_i\| \leq b_i d_i$, and $\|u_i\| \leq b_i(a_i + b_i)d_i$ almost everywhere. If $c_i \rightarrow c_i^\infty$, then $x_i \rightarrow c_i^\infty$, $v_i \rightarrow 0$.

Proof. Equation (36) factors as $\dot{y}_i = a_i(c_i - y_i)$, $\dot{x}_i = b_i(y_i - x_i)$. Variation of constants expresses each state as its initial state and previous-stage input with nonnegative weights of total one. Closed convexity proves invariance, including bounded measurable inputs. The identities $v_i = b_i(y_i - x_i)$ and $u_i = a_i b_i(c_i - y_i) + b_i^2(x_i - y_i)$ give the bounds. For convergence, split each exponentially weighted convolution at a sufficiently late time. Its finite early part decays and its remaining input error is arbitrarily small. Apply this first to y_i , then to x_i . $\square$

Assign fixed offsets r_i and anchors in one compact convex K . With $C_i = r_i + K$, choose $c_i(t) = r_i + w_i(t)Z$, interpreting each local row on its applied support. A locally finite action embedding tending to infinite physical time and Theorem 2 imply

$$x_i(t) - r_i \rightarrow \mu_\sigma Z, \quad v_i(t) \rightarrow 0. \quad (37)$$

For a common paired future, restoration is exactly $\pi_\sigma^\top(Y(q) - Y_{\text{ref}})Z = 0$. For Theorem 1's round trip it is $(\Theta q - g)Z = 0$. Theorem 6 also applies to the reference-assisted law (34). With point agents, all $r_i = 0$ gives common-point rendezvous. Positive bodies with separation bounded below cannot all converge to one point; separated offsets give a common formation translation instead.

B. Exact held-goal evolution

The continuous controller has an exact discrete observation map whenever c_i is held on an interval of length $h > 0$.

Suppress the body index and put $A_h = e^{-ah}$, $B_h = e^{-bh}$. Then

$$\begin{aligned} y^+ &= A_h y + (1 - A_h)c, \\ x^+ &= B_h x + K_h y + (1 - B_h - K_h)c, \\ K_h &= \begin{cases} b(B_h - A_h)/(a - b), & a \neq b, \\ b h e^{-bh}, & a = b. \end{cases} \end{aligned} \quad (38)$$

Here $K_h = b \int_0^h e^{-b(h-s)} e^{-as} ds$. Its integral representation proves $K_h \geq 0$ and $K_h \leq 1 - B_h$, so all displayed weights are nonnegative and sum to one. This proves the sampled states remain in the same convex cells in exact arithmetic. Near $a = b$, a stable evaluation or the limiting formula is needed; a subtractive exponential formula is not a numerical certificate. The trajectory checks below use fixed separated gains or the equal-gain branch. They do not treat forward Euler as an exact implementation of the controller.

The initial constraint on $y = x + v/b$ is substantive. In one dimension, take $C = [-1, 1]$, $x(0) = 1$ and $v(0) > 0$. Continuity forces $x(t) > 1$ for small positive t under any locally bounded acceleration, irrespective of the goal. Position membership alone cannot guarantee cell invariance.

C. A finite-workspace construction retaining every body

Bodies are closed balls of radius $r_b > 0$. Fix $R > 0$, $\Delta > 2r_b$, and finitely many horizontal bases $r_\ell = (p_\ell, 0)$ with

$$\|p_\ell - p_k\| \geq 2R + \Delta \quad (\ell \neq k). \quad (39)$$

Cell $C_\ell = r_\ell + \bar{B}_R(0)$ belongs to a column. Each body has an immutable assigned column and is used for at most one active incarnation. All bodies exist initially: one may be active in each column, and the others wait at rest at $(p_\ell, H + j\Delta)$, $j = 0, 1, \dots$, with $H = R + \Delta$. There are no other bodies or obstacles in the reserved cells, queue paths, and parking paths. Initially active bodies satisfy $x, y \in C_\ell$. All objective anchors lie in $\bar{B}_R(0)$.

A column controller serializes its home, brake, park, clearance, entry, and queue-advance phases. At most one active, homing, braking, parking, or entering body holds its transition reservation. It accepts only a finite request list with available unused queue bodies, sufficient parking capacity, and nonempty logical support at every committed record prefix. A request that removes the sole logical participant before a permitted birth is not admissible. Different columns may progress concurrently. The construction assumes a correct sequencer and a correct reservation controller; it does not establish decentralized fault-tolerant agreement.

Allocate M_ℓ parking berths, at least the number of requested retirements and no more than the number of bodies assigned to column ℓ . Such an allocation exists because each body is used at most once. The j th retiree receives

$$z_{\ell,j} = -R - (M_\ell - j + 1)\Delta, \quad 1 \leq j \leq M_\ell. \quad (40)$$

Earlier retirees occupy deeper berths. A parked body remains there at rest and is never deleted or reused. The reserved centre workspace is finite: its lowest height is $-R - M_{\max}\Delta$, and its highest is the greatest initial queue height. Body clearance additionally expands the reserved workspace by r_b .

Require $A_{\max}, V_{\max} > 0$ and $2b_i R \leq V_{\max}, 2b_i(a_i + b_i)R \leq A_{\max}$. The plant is the ideal double integrator of Theorem 6; thrust allocation, attitude dynamics, gravity compensation, sensor error, and contact mechanics are not supplied by that model. Selected active data links must satisfy $\|r_i - r_j\| + 2R \leq R_{\text{comm}}$. Lifecycle reports use a separately specified reliable channel covering the whole reserved workspace, including the parking berths; active data-radio range alone does not establish that coverage. Every report binds a column and the expected incarceration.

Home and brake. At local retirement, priority operations stop and the goal becomes r_i . Choose $m \in (0, R)$, $0 < A_s \leq A_{\max}$, and $v_* > 0$ with $v_* \leq V_{\max}$ and $v_*^2/(2A_s) \leq m$. There is a fixed duration T_h such that every admissible home state reaches $\|x - r_i\| \leq R - m$, $\|v\| \leq v_*$. Indeed, with

$$\phi_i(t) = e^{-b_i t} + b_i \int_0^t e^{-b_i(t-s)} e^{-a_i s} ds, \quad (41)$$

variation of constants gives $\|x - r_i\| \leq R\phi_i(t)$ and $\|v\| \leq b_i R(e^{-a_i t} + \phi_i(t))$. Both bounds tend to zero. A finite pool permits a common maximum T_h . Brake for the fixed time $T_s = v_*/A_s$ with constant acceleration $-v_0/T_s$. The path moves at most $v_*^2/(2A_s) \leq m$, remains in the cell, respects the actuator bounds, and ends at rest. No physical state is reset.

Finite parking. Hold the resulting horizontal coordinates. Let $z_b \in [-R, R]$ be the braking endpoint and $L_i = z_b - z_{\ell,j} > 0$. Set $L_{\max} = 2R + M_{\max}\Delta$, choose $0 < V_e \leq V_{\max}$, and use the fixed duration

$$T_p \geq \max\{2\sqrt{L_{\max}/A_{\max}}, 2L_{\max}/V_e\}. \quad (42)$$

For normalized time $u \in [0, 1]$, let $s(u) = 2u^2$ on $[0, 1/2]$ and $s(u) = 1 - 2(1 - u)^2$ on $[1/2, 1]$. The downward path is

$$x(t) = x_b - L_i s(t/T_p) e_3, \quad 0 \leq t \leq T_p. \quad (43)$$

It is rest-to-rest and monotone in height. Its peak speed is $2L_i/T_p \leq V_e$, and its acceleration magnitude is $4L_i/T_p^2 \leq A_{\max}$ almost everywhere. After reaching the berth, hold at rest and send a matching parked-clearance report. Its truth persists through delayed receipt. This replaces the unbounded permanent cruise in v0.2.

Entry and queue advance. After the clearance is received, the front queued body moves vertically from H to home, rest-to-rest, using the same normalized profile with duration $T_{\text{in}} \geq \max\{2\sqrt{H/A_{\max}}, 2H/V_e\}$. Only then is its birth record emitted. It waits at home without a priority row until coherent bootstrap, starts Theorem 6 from $x = y = r_i$, or follows the retirement procedure without sending when removal is already known. Remaining queue bodies descend one level, lowest first; each finishes a rest-to-rest move before the next starts, with $T_q \geq \max\{2\sqrt{\Delta/A_{\max}}, 2\Delta/V_e\}$. Queue advance finishes before the next entry. All phase durations can be fixed independently of q ; common-future comparisons also require common admissible control-message timing.

Theorem 7 (finite-workspace staged safety). Under the stated initial, reservation, channel, gain, capacity, and request-admissibility conditions, every prescribed replacement finishes in finite time. Every body remains present, respects the speed and acceleration bounds, and has centre separation at least Δ

from every other body at all times. Retirees eventually remain at rest in finite parking space.

Proof. Theorem 6 gives the active and home bounds; the braking margin proves cell containment and the braking actuator bounds. The normalized rest-to-rest paths prove the entry, parking, and queue bounds. Position and velocity are continuous at every phase switch. Acceleration jumps are allowed; no jerk bound is claimed.

Every body stays within the horizontal radius- R footprint of its assigned column. Equation (39) therefore gives horizontal separation at least Δ between different columns, at every height.

Within one column, a cell occupant has height between $-R$ and R , while every waiting queue body has height at least H , giving a vertical gap of at least Δ . Entry descends away from upper waiting bodies. In queue advance, the lower body moves away first, and the next moving upper body stops at least Δ above the lower body's already reached height. Thus the queue ordering preserves separation.

A j th parking path never descends below $z_{\ell,j}$. Every earlier parked body has height at most $z_{\ell,j} - \Delta$, so parking never approaches an older retiree closer than Δ . After parking, the newest berth has height at most $-R - \Delta$; every later active or braking body remains above $-R$, and every later entering body remains at nonnegative height until it starts the cell controller. The reservation prevents overlapping entry and retirement in the same cell. Induction over the finite request list covers all pairs, including all retired bodies. Since $\Delta > 2r_b$, the bodies are disjoint.

Each motion phase has finite prescribed duration. Available queue bodies, sufficient berths, the serialized finite request list, bounded control-report receipt, and bounded coherent bootstrap/application ensure that every phase enabling condition is eventually reached. No physical disappearance or unbounded cruise is used. $\square$

For an explicit instance, take $R = \Delta = V_{\max} = A_{\max} = 1$, $r_b = 1/10$, $a_i = b_i = 1/10$, $m = 1/4$, $v_* = 1/5$, and $A_s = V_e = 1$. A home duration $T_h = 10$ is sufficient: $\phi_i(10) = 2/e < 3/4$ and the speed bound is $3/(10e) < 1/5$. Braking takes $1/5$ second and moves at most $1/50$. With $M_{\max} = 3$, use $T_p = 10$, $T_{\text{in}} = 4$, and $T_q = 2$. Berths in a three-capacity column are at heights $-4, -3, -2$. Base spacing three gives geometric separation; an active edge between adjacent bases needs radio range at least five. Control-channel coverage of lower berths remains a separate requirement.

The construction consumes $O(N_p)$ body and berth records, $O(N_p)$ physical queue/parking capacity, reserved disjoint footprints, a correct reservation/sequencing service, and the stated control-channel coverage. It does not handle arbitrary approaching trajectories, cluttered space, perpetual churn, or reuse of parked bodies without a new construction.

D. From physical service to action contracts

Suppose every physical interval of length h contains at most $C + \lceil \rho h \rceil$ primitive actions, including simultaneous bursts. Define $N(h) = C + \lceil \rho h \rceil + 2$. Complete records within p seconds of their emission and complete newborn bootstrap or terminal disposition within p seconds of birth-record emission.

Deliver data within d , and serve/supersede ready copies within q_s . Then Lemma 2 holds with $P = N(p)$, $D_{\text{net}} = N(d)$, $Q = N(q_s)$; physical use age is at most $\max(p, d) + q_s$.

On a fixed final strongly connected graph, send on each edge and activate each receiver at least once per τ seconds, using FIFO receive weights in $[\eta, 1 - \eta]$, $0 < \eta \leq 1/2$. A source waits behind at most $n - 2$ others, so service takes at most $(n - 1)\tau$. Each edge is actually consumed within $n\tau + d$ seconds after any phase start. Hence $B = N(n\tau + d) + 1$ consecutive action indices contain a strongly connected actual influence union: any shorter physical span would violate the counting bound.

For a finite pool N_p , let $E_p = N_p(N_p - 1)$. Assume fixed-phase periodic timers, with no extra reset-triggered firings: exactly one scheduled send per directed edge per period τ , at most one priority activation per node per period τ , one delivery per send with delay at most d , and at most $M_{\log}$ additional record/bootstrap primitives in the entire finite change interval. With the corresponding at-least-once service convention for active timers, one sufficient choice is $\varrho = (2E_p + N_p)/\tau$, $C = \lceil 2E_p + N_p + E_p d/\tau \rceil + M_{\log} + 2$, and $q_s = (N_p - 1)\tau$. Waiting receivers have no ready FIFO before bootstrap; retiring receivers terminally supersede their buffers. These bounds include those lifecycle dispositions. Geometry alone proves neither timely delivery nor processing. Sequence and reception counters may grow with execution length; a bound on retained records is not a bound on their bit representation. Timing and availability must also be independent of q when using Theorem 3's paired comparisons.

Combining Theorems 2, 6 and 7 yields safe formation convergence after finite changes for the specified staged nominal model. It does not prove arbitrary outside approach, obstacle avoidance, six-degree-of-freedom flight, noise/fault robustness, perpetual-churn convergence or hardware performance.

VII. EXECUTABLE VERIFICATION AND EVIDENCE BOUNDARIES

A. Fresh checks, not inherited counts

The supplied v0.2 text reports historical A01–I04 suites without supplying their code, logs, or claim maps. Those counts are not promoted to reproduced evidence. This revision supplies `verification/verify.py`, a fixed seed (20261003), machine-readable results, and raw CSV trajectories. Exact algebra uses Python rational arithmetic. Trajectory evaluation uses NumPy float64 and the analytic held-input maps, not a generic ODE integrator. The source hash and software versions appear in `verification_report.json`. A scenario means one parameterized fixture; it may contain many assertions. The complete run has 260 scenarios.

The 72 round trips include 46 cases with a strictly positive restoring q_* , 12 zero-influence immediate returns, and 14 cases whose algebraic q_* is outside the positive simplex. A separate zero-influence fixture repeats one incumbent row into both survivors; it has $\Theta = 0$ but $g = (-1/4, 1/4) \neq 0$, so changing the newcomer cannot restore the reference. Exact comparison evaluates both the direct join–mix–condition history and the coefficient formula. The 32 variable-delay tests exhaust the two

TABLE I
NEW FINITE CHECKS. NO ROW IS A UNIVERSAL PROOF.

Fixture family	Cases
Fixed-delay invariants and common columns	54
Fixed-history round trips	72
Variable-delay common-column fixture	32
Packet timing traces; receiver FIFO unit cases	12; 4
Historical skew families; quantizer boundaries	40; 11
Nonpolynomial bootstrap algebra	3
Exact counterexamples; tensor/boundary fixtures	17; 4
Information-latency seams	3
Reference injection; bootstrap clock; zero gain	3; 1; 1
Held-goal trajectories; finite parking trajectory	2; 1

activation orders and all binary lag choices within one specified two-node, four-action fixture; they do not exhaust admissible networks or infinite schedules. Receiver readiness cases are unit tests, not end-to-end distributed protocol certification. Three exact reference-assisted runs check the augmented error bound for three nodes, delay depth two, one activation per node per three actions, injection weight $1/4$, and 60 actions. The bootstrap-clock counterexample tests the clarified deadline origin; it is not included among compliant post-repair executions.

B. Exact restoring and obstructed histories

For a three-incumbent example use

$$W_0 = \begin{bmatrix} 3/5 & 1/5 & 1/5 \\ 1/5 & 1/2 & 3/10 \\ 3/10 & 1/5 & 1/2 \end{bmatrix}, \quad (44)$$

constant initialized histories, final path 1–2–3, enlarged path 1–2–3–4, $c = 1/4$, $\beta_i = 1/4$, $\gamma = 3/4$, and a stay of four updates. With uniform unit lag on every directed edge, the reference is $(57/160, 5/16, 53/160)$, $\Theta = 37/256$, and

$$q_* = (99/370, 119/370, 76/185). \quad (45)$$

The exact direct returned endpoint equals that reference.

For a directional instance, list lags by (receiver, source). Final lags $(1, 2)$, $(2, 1)$, $(2, 3)$, $(3, 2)$ are $(0, 1, 2, 0)$; enlarged lags $(1, 2)$, $(2, 1)$, $(2, 3)$, $(3, 2)$, $(3, 4)$, $(4, 3)$ are $(1, 0, 2, 1, 0, 2)$. Then the reference is $(28/75, 7/25, 26/75)$, $\Theta = 1/8$, and $q_* = (27/50, 17/400, 167/400)$. The exact direct endpoint again equals the reference. In contrast, the actual packet futures in Section IV require incompatible returned states $11/16$ and $7/8$ to restore the same prescribed first coordinate. These are different claims: fixed-history solvability does not imply one bootstrap works for both futures.

C. Controller trajectories

Two separate four-body checks use a ring, $c = 1/4$, constant initialized histories, and

$$W_0 = \frac{1}{10} \begin{bmatrix} 4 & 3 & 2 & 1 \\ 1 & 4 & 3 & 2 \\ 2 & 1 & 4 & 3 \\ 3 & 2 & 1 & 4 \end{bmatrix}. \quad (46)$$




Fig. 1. Augmented priority error in two fixed-support delay cases. Very small values reach floating-point resolution; displayed values below 10^{-16} are floored for the logarithmic axis only. Raw data are unchanged. This is not a finite-time consensus theorem.




Fig. 2. Position tracking under the exact held-goal kernel. This evaluates the ideal double-integrator construction, not the flying-cat rigid-body controller.

The anchors are the four vectors $(\pm 1/2, \pm 1/2, \pm 1/2)$ with an even number of negative coordinates. The offsets are $(0, 0, 0)$, $(3, 0, 0)$, $(3, 3, 0)$, $(0, 3, 0)$. Initial $x - r$ is 0.7 times the reversed anchor list and initial $y - r$ is 0.6 times the anchor list. Thus both belong to radius-one cells. Controller gains are $a = 0.7$, $b = 0.5$. Priority updates run at 12 Hz and the held-input controller is evaluated at 120 Hz for 60 seconds.

Uniform lags are two. Directional lags on $(1, 2)$, $(2, 1)$, $(2, 3)$, $(3, 2)$, $(3, 4)$, $(4, 3)$, $(1, 4)$, $(4, 1)$ are $(0, 2, 1, 3, 2, 0, 3, 1)$. The analytic endpoints are, respectively, $(1/4, 1/4, 1/4, 1/4)$ and $(9/35, 9/35, 8/35, 9/35)$. At 60 seconds the evaluated augmented priority errors are at float64 roundoff level; position errors are below 4×10^{-13} . The two runs have sampled minimum centre separation 2.40416 m, maximum speed 0.46030 m/s, and maximum acceleration 0.36118 m/s². These sampled values are not global optimizers of the trajectory extrema. The analytic bounds remain separation at least 1 m, speed at most 1 m/s, and acceleration at most 1.2 m/s².

D. Retained-body lifecycle trajectory

The finite-parking fixture has five physical bodies: an active body and two queued bodies in one column, and one stationary active body in each of two other columns. It executes two replacements with the numerical parameters of Theorem 7, a clearance-message delay of 0.3 seconds, and a five-second active interval before the second retirement. The first two retirees end at rest at heights -4 and -3 . All five bodies remain in the state through 121 seconds. At 12,102 sampled times,




Fig. 3. Sampled all-body separation across two staged replacements, retaining both parked bodies. Phase construction supplies the continuous-time argument.

including phase endpoints, the minimum centre separation is 1 m, the maximum speed is 1 m/s, and the maximum acceleration is 1 m/s². The theorem, not this sampled minimum, establishes continuous-time safety. Raw positions, velocities, accelerations, and phase boundaries are supplied. At an acceleration switch the CSV records one branch value; the analytic bounds apply to both one-sided limits. It does not sample or bound jerk.

E. What these artifacts do not establish

The packet fixtures, fixed-history experiments, controller trajectories, and parking trajectory test separate contracts. They do not certify one integrated real-time packet-and-flight implementation. The supplied `Flying_Cats(1).html` uses a 120 Hz rigid-body engine with thrust/torque and predictive avoidance; its removal handler deletes a rigid body, and its restoration trial is restricted to zero delay. Those are source-level correspondence facts, not an assessment of its visual quality. Its trajectories cannot be substituted for evidence of Theorem 7 or the membership-prefix protocol.

No hardware run, stochastic population-level confidence claim, proof-assistant certificate, or independent human acceptance is asserted. Every displayed value is generated by the new package, with raw values retained before display rounding. The separate audit records the exact supplied input hashes and the version-by-version correction ledger.

VIII. CONCLUSION

The complete delayed history determines the fixed-delay endpoint. Proportional departure yields an exact restoration equation whose three-dimensional projection can have more solutions than priority restoration. Actual packet history and future scheduling can obstruct a common restoring bootstrap. Smooth-state and finite-bit lower bounds identify historical information required for the declared survivor-query interface, while a counted reference service supplies eventual restoration under a modified law. The continuous filter transfers priority convergence to point rendezvous or separated formation convergence. Physical entry and retirement are feasible under the finite-workspace staged construction, including every parked retiree, with explicit berth capacity and control-channel coverage. The information, service and geometry hypotheses delimit these results.

A. Retained scope and resource limits

The proofs retain exact real arithmetic, positive probability data, correct identities and metadata, bounded used ages, and the graph/service assumptions stated for each result. Fixed-delay restoration uses globally supplied histories and prescribed mixing. Universal-future equations assume one parameter-independent admissible family. Historical memory bounds are for the declared erasure and query interface, with neither deleted values nor input-dependent side channels given for free. Safe formation uses invariant cells, finite single-use bodies and berths, a correct controller, and reliable lifecycle reports.

Remaining research questions include finite-bit causal maintenance through turnover, noise-robust physical control, efficient computation of unknown-future certificates, and an end-to-end implementation of the combined contracts. They are not missing premises concealed in the retained theorems. No claim is made for arbitrary unannounced faults, Byzantine agents, unrestricted outside entry, obstacle-rich environments, indefinite body reuse, or perpetual-churn convergence.

AI DISCLOSURE AND RESPONSIBILITY

This work was developed with assistance from AI agents operating within the HORTAME 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.

APPENDIX A

AN INFORMATION-LATENCY BENCHMARK

Let $\delta \geq 0$. This section concerns an observer with no informative channel about a fresh exogenous $q \in \Delta_n^+$ during $[0, \delta)$. Prior memory and independent randomness do not contain this new input. The prescribed immediate joined-row mean has the form $\nu(q) = a + \kappa(q, 0)$, $\kappa = \gamma/(n+1)$, with known a . Its centre is $c = \nu(\mathbf{1}_n^\top/n)$, and $\rho^2 = \kappa^2(1 - 1/n)$.

The minimax squared error before disclosure is exactly ρ^2 ; the hidden-interval integral is exactly $\delta\rho^2$. For proof, use positive near-vertex rows $q^j = \epsilon\mathbf{1}_n^\top + (1 - n\epsilon)e_j^\top$. Their targets have mean c and radius $(1 - n\epsilon)\rho$. For any estimate x , the average squared error is $\|x - c\|^2 + (1 - n\epsilon)^2\rho^2$. Its law is the same under all hidden choices. Taking expectations and then $\epsilon \downarrow 0$ proves the pointwise bound. Integrating the same nonnegative identity proves the integral bound with one fixed hidden input over the whole interval. Holding c attains both suprema. This is a prescribed-target information result, not a mean-conservation claim for asynchronous dynamics.

For a separate physical benchmark, let $2 \leq n \leq 4$ and let $z_1, \dots, z_n$ be an affinely independent regular simplex in $\mathbb{R}^3$ centred at zero with radius R_s , and let the hidden goal be

$g(q) = \kappa qZ$, of supremal radius $r = \kappa R_s$. An isolated double integrator starts at rest at zero, stays in the goal hull, and has speed/acceleration bounds $V, A > 0$. Exact goal disclosure occurs at δ ; impose terminal rest at horizon s . Define

$$d_{A,V}(h) = \begin{cases} Ah^2/4, & h \leq 2V/A, \\ Vh - V^2/A, & h \geq 2V/A. \end{cases} \quad (47)$$

The deterministic minimax endpoint error is r for $s \leq \delta$ and $\max(0, r - d_{A,V}(s - \delta))$ otherwise. Indeed terminal rest implies $\|v(t)\| \leq \min(At, A(h - t), V)$; integrating yields the rest-to-rest distance (47), attained radially by accelerate/decelerate or accelerate/cruise/decelerate motion. Average any policy over the finite orthogonal simplex-permutation group on positive near-vertex goals. Convexity preserves dynamics, constraints and terminal rest and cannot increase worst endpoint error. Before disclosure its averaged state is zero, since the group's fixed subspace in the affine span of the regular simplex is zero. The hull constraint places the actual state in that span, including segment and triangle embeddings. The distance bound therefore gives the lower error after disclosure. Holding zero, then moving radially toward the disclosed goal through distance $\min(\|g\|, d_{A,V})$, attains it while staying in the hull. For the tetrahedron $(\pm 1/2, \pm 1/2, \pm 1/2)$ with an even number of negative coordinates, $R_s = \sqrt{3}/2$ and this is a three-dimensional benchmark. It is not an optimality theorem for the networked formation construction.

REFERENCES

- [1] G. Bolduc, "Priority-State Dynamics in Open Multi-Agent Systems: Dynamic Membership and History-Dependent Consensus Endpoints," full edition, version 07, supplied preprint, Sep. 2026.
- [2] S. Naderi, M. J. Blondin, and S. Roy, "Consensus in Multi-Agent Systems with Uniform and Nonuniform Communication Delays," arXiv:2603.16523v1, 2026. [Online]. Available: <https://arxiv.org/abs/2603.16523v1>.
- [3] K. I. Tsianos and M. G. Rabbat, "The Impact of Communication Delays on Distributed Consensus Algorithms," arXiv:1207.5839v1, 2012. [Online]. Available: <https://arxiv.org/abs/1207.5839v1>.
- [4] A. V. Proskurnikov and G. C. Calafiore, "Delay Robustness of Consensus Algorithms: Beyond the Uniform Connectivity (Extended Version)," arXiv:2105.07183v1, 2021. [Online]. Available: <https://arxiv.org/abs/2105.07183v1>.
- [5] D. Deplano, N. Bastianello, M. Franceschelli, and K. H. Johansson, "Optimization and Learning in Open Multi-Agent Systems," arXiv:2501.16847v1, 2025. [Online]. Available: <https://arxiv.org/abs/2501.16847v1>.
- [6] Y. Cherapanamjeri, S. Garg, N. Rajaraman, A. Sekhari, and A. Shetty, "The Space Complexity of Learning-Unlearning Algorithms," arXiv:2506.13048v1, 2025. [Online]. Available: <https://arxiv.org/abs/2506.13048v1>.
- [7] P. Sekercioglu, N. De Carli, and D. V. Dimarogonas, "Safe Formation Control of Open Multi-Robot Systems with Connectivity-Preserving Reconfiguration," arXiv:2609.29939v1, 2026. [Online]. Available: <https://arxiv.org/abs/2609.29939v1>.
- [8] J. Hu, K. H. Johansson, and A. I. Rikos, "Distributed Coordination Algorithms with Efficient Communication for Open Multi-Agent Systems with Dynamic Communication Links and Processing Delays," arXiv:2603.08038v1, 2026. [Online]. Available: <https://arxiv.org/abs/2603.08038v1>.
- [9] The Stacks Project Authors, "Noetherian rings," Lemma 10.31.1, tag 00FM. [Online]. Available: <https://stacks.math.columbia.edu/tag/00FM>. Accessed: Oct. 3, 2026.