

Priority-State Dynamics in Open Multi-Agent Systems: Dynamic Membership and History-Dependent Consensus Endpoints

Guillaume Bolduc

Abstract—We characterize the mean priority returned by a single-agent arrival–mixing–departure history. Each row is a probability vector over the active objectives, so removing an agent also removes a coordinate and conditions the surviving rows. For arbitrary nonnegative row-stochastic mixing, we derive the returned mean directly from the incumbent state, retention masses and mixing matrix. Holding those inputs fixed gives the unique restoring newcomer conditional priority when its influence is positive, an exact strict-positivity feasibility test, and a separate zero-influence case. Under equal retention and doubly stochastic mixing, restoration requires a weighted average of incumbent rows; for one unweighted consensus update, this is the newcomer-neighbor row mean. The general result is an offline input-level characterization, not a distributed arrival protocol. An exact counterexample excludes unrestricted endpoint invariance. Standard final-epoch average consensus supplies the limiting-priority interpretation; no new physical-rendezvous or persistently open convergence theorem is claimed.

Index Terms—priority consensus, open multi-agent systems, membership changes, stochastic matrices, simplex, multi-objective optimization.

I. QUESTION AND CLOSEST PRIOR WORK

When each participant owns an objective, membership changes affect both the owners of priority rows and the coordinates those rows describe. A probability-preserving dimension change need not preserve the eventual agreed priority. We ask when adding one agent, mixing, and removing that same agent restores the original population’s priority endpoint. The setting is motivated by prioritized optimization [1], [2]; the result concerns the probability-state transformations themselves.

Naderi and Blondin study changing priority values for a fixed agent/objective set and a fixed-dimensional priority matrix [2]. The later delay analysis of Naderi, Blondin and Roy also retains fixed membership while allowing stale neighbor information [3]. Here membership changes the owners of rows and the objective coordinates within every surviving row. Priority-value variation, delayed information and membership-state transformation are distinct questions. The results below do not establish a combined delayed-and-open-membership algorithm.

The closest aggregate-preservation comparison is Oliva *et al.* [4]. Their open-system framework maintains neighbor

storage that cancels departed agents’ accumulated contributions; the maintained aggregate equals the active agents’ last-activation values. Here departure conditions probability rows after removing an objective coordinate. No auxiliary flow memory cancels a departing agent’s influence, so the aggregate generally changes. We characterize equality for the prescribed maps and a prescribed round trip, rather than supplying an invariant-enforcing protocol.

Zhou *et al.* explicitly address changing global state dimension: their open state vector removes departed agent blocks, retains surviving blocks and appends arrivals [5]. Each individual state remains in a fixed space $\mathbb{R}^d$; membership changes the number of blocks in the global stack. In contrast, the coordinates of each priority row here index active objectives, so departure deletes a coordinate from every survivor and requires a rule for its probability mass. Their weights are coefficients in a population aggregate of agent states; the entries of w_i here are agent i ’s probability allocation over objective identities. The distinction is semantic, not the choice of the words weights or priorities. Zhou *et al.* study bounded dynamic weighted-average consensus under sampled-data impulsive communication; the present result characterizes exact mean restoration for a prescribed finite round trip. Thus changing state dimension is not itself a novelty claim. Recent work by Jiang *et al.* uses auxiliary and surplus variables for sum-of-values conservation in directed open double-integrator systems [6]. Its protocol addresses historical influence and directed imbalance, whereas the matrices below encode mixing followed by conditioning. These different state and update rules must be distinguished when comparing endpoint claims.

Open-system modeling, disagreement and stability have a substantial lineage [7]–[10]. Open optimization and operator frameworks address broader continuing changes [11], [12]. Deplano *et al.* formulate open operators acting between labeled spaces of changing dimension and composition [12]; neither variable dimension nor state maps in general are contributions here. The present object is the particular active-objective probability semantics and its prescribed-history endpoint characterization. Open average/ratio-consensus protocols add accounting to recover specified aggregates [13], [14]; switching-consensus analyses impose their own graph and timing hypotheses [15]. None is replaced by the finite-history calculation here. Proportional deletion itself is standard subcomposition and closure [16].

Hu, Johansson and Rikos use quantized mass/count states with arrival initialization and departure handoff to preserve averaging sums on dynamic directed networks [17]. Their processing-delay extension restricts transmissions to recipients that remain active through the delay horizon. Those accounting rules differ from deleting an objective coordinate and conditioning every surviving priority row. The result here concerns restoration for the specified probability-state maps; the browser's synchronized transport barrier is not their protocol.

Our focus is the offline input-level characterization in Theorem 4.1: with the incumbent state, retention masses and mixing fixed, it determines whether a restoring newcomer row exists and what it must be. Theorem hypotheses require only a nonnegative row-stochastic mixing matrix. Double stochasticity is added in Corollary 4.2; its one-update case has a local neighbor-row interpretation. Neither a distributed online arrival protocol nor robustness to an unknown departure time is claimed. An accompanying browser simulation exposes the priority maps and a prescribed equal-retention restoration trial [18]; its physical target is computed from the epoch mean, independently of the evolving local rows.

II. PRIORITY STATE AND ORDINARY EVOLUTION

For $n \geq 1$, define

$$\begin{aligned} \Delta_n &= \{u \in \mathbb{R}^n : u_j \geq 0, \mathbf{1}_n^\top u = 1\}, \\ \Delta_n^+ &= \{u \in \Delta_n : u_j > 0 \ \forall j\}, \end{aligned} \quad (1)$$

$$\begin{aligned} \mathcal{S}_n &= \{W \in \mathbb{R}^{n \times n} : W \geq 0, W\mathbf{1}_n = \mathbf{1}_n\}, \\ \mathcal{S}_n^+ &= \{W \in \mathcal{S}_n : W_{ij} > 0 \ \forall i, j\}. \end{aligned} \quad (2)$$

Here $W \geq 0$ is entrywise. In particular, $\mathcal{S}_1^+ = \{[1]\}$; “positive” does not impose an upper bound strictly below one in the singleton case. Row $w_i^\top$ records agent i 's priorities over the ordered active objectives. Rows and columns share labels but represent different roles: the row owner and the objective owner.

The mean priority is the row vector

$$m(W) = \frac{1}{n} \mathbf{1}_n^\top W. \quad (3)$$

It averages current rows, not necessarily original declared preferences. The distinction becomes essential after events and in systems with memory.

A. Connection to a consensus endpoint

For $n \geq 2$, let G be a connected simple undirected graph, H its zero-diagonal adjacency matrix, $d_i = \sum_j H_{ij}$, $d_{\max} = \max_i d_i$, and $L = \text{diag}(d) - H$. Choose

$$0 < c < d_{\max}^{-1}, \quad P = I - cL, \quad W(k+1) = PW(k). \quad (4)$$

The updates in (4) are synchronous mathematical iterations: every row uses the corresponding common previous state. They do not assert that real packets have zero age. For $n = 1$, define $P = [1]$ and omit the gain test.

Lemma 2.1 (Admissibility, mean and spectral contraction). Under (4), P is nonnegative, symmetric and doubly stochastic with positive diagonal. It maps $\mathcal{S}_n$ into itself and $\mathcal{S}_n^+$ into

itself. If $n \geq 2$, there exists $0 \leq \rho < 1$ such that, for every $W \in \mathbb{R}^{n \times n}$ and $k \geq 1$,

$$\|P^k W - \mathbf{1}_n m(W)\|_F \leq \rho^k \|W - \mathbf{1}_n m(W)\|_F. \quad (5)$$

For $k = 0$ the same inequality holds with factor one. Moreover $m(PW) = m(W)$. For $n = 1$, consensus is already attained.

This is the standard fixed-epoch average-consensus result [19]. Its self-contained proof is supplied in Appendix A, together with the storage-coordinate qualifications. Consequently, a returned state U and reference state W on the same ordered population have the same eventual priority under (4) exactly when $m(U) = m(W)$. No equality of intermediate means is required. The mixing during the newcomer stay will be allowed to be more general than (4); the endpoint interpretation uses a subsequent fixed average-consensus epoch.

III. PROBABILITY-STATE MEMBERSHIP MAPS

The following maps apply to positive priority states. Strict positivity ensures that every nonempty survivor set has positive retained mass. An empty population is outside this state space. These are specifications of state semantics, not certification of a saved distributed implementation.

A. Departure

Let $\emptyset \neq A \subseteq V$ be the ordered survivor set and $m = |A|$. Define

$$\begin{aligned} Q_A(W) &= S_A W S_A^\top, \quad s_A(W) = Q_A(W) \mathbf{1}_m, \\ \mathcal{D}_A(W) &= \text{diag}(s_A(W))^{-1} Q_A(W). \end{aligned} \quad (6)$$

For $W \in \mathcal{S}_n$, this map is defined exactly when every retained row has strictly positive mass on surviving objectives. If a retained row puts all its mass on removed objectives, (6) is undefined; selecting a replacement row would be a different, explicitly specified rule.

Proposition 3.1 (Departure and uniqueness). For $W \in \mathcal{S}_n^+$ and nonempty A , $\mathcal{D}_A(W) \in \mathcal{S}_m^+$. It is the unique stochastic transformation of each surviving row that preserves all ratios between its surviving coordinates. More generally it belongs to $\mathcal{S}_m$ whenever $W \in \mathcal{S}_n$ and all entries of $s_A(W)$ are positive. *Proof.* For retained row i , let $s_i = \sum_{j \in A} w_{ij}$. Its new entries are w_{ij}/s_i . They are nonnegative, sum to one, and are positive if the original entries were positive. In the positive case the common factor $1/s_i$ preserves every ratio. Conversely, ratio preservation forces all transformed coordinates to be $a_i w_{ij}$ for one scalar a_i ; summing them forces $a_i = 1/s_i$. $\square$

This is conditioning a positive composition on its retained coordinates, rather than inventing a new normalization primitive [16]. It preserves each surviving agent's relative priorities, not the prior population's aggregate mean.

B. Arrival

Arrival is a parameterized family of admissible ratio-preserving maps. Append a new label. Existing row i assigns it a chosen mass $\beta_i \in (0, 1)$, and the newcomer supplies a chosen $v \in \Delta_{n+1}^+$. Define

$$\mathcal{J}_{\beta, v}(W) = \begin{bmatrix} \text{diag}(\mathbf{1}_n - \beta) W & \beta \\ v_{1:n}^\top & v_{n+1} \end{bmatrix}. \quad (7)$$

Neither β nor v is selected by the admissibility proof. Uniqueness applies only to the incumbent old-coordinate scaling once its insertion mass is fixed; there is no unique canonical arrival rule without an additional preference policy.

Proposition 3.2 (Arrival and immediate inverse). The map (7) sends $\mathcal{S}_n^+$ into $\mathcal{S}_{n+1}^+$. For fixed β_i , its old-coordinate portion is the unique ratio-preserving part of the updated existing row. If the appended agent is immediately removed, with no intervening evolution or other event, then

$$\mathcal{D}_V(\mathcal{J}_{\beta,v}(W)) = W. \quad (8)$$

Proof. Existing row i has positive entries summing to $(1 - \beta_i) + \beta_i = 1$; the new row is positive stochastic by hypothesis. Ratio preservation multiplies all old coordinates by a common scalar, necessarily $1 - \beta_i$. Immediate departure removes the mass β_i and divides the remaining row by $1 - \beta_i$, recovering $w_i^\top$; the newcomer's row is removed. $\square$

Several arrivals can be represented as a specified ordered composition of these maps. The order and the newcomers' declarations are part of the model. Compositions preserve admissibility, but distinct orders need not preserve the same values.

IV. ENDPOINT RESTORATION AFTER A ROUND TRIP

A. A join-evolve-depart counterexample

Start from

$$W = \begin{bmatrix} 9/10 & 1/10 \\ 1/10 & 9/10 \end{bmatrix}, \quad \beta = \begin{bmatrix} 1/3 \\ 1/3 \end{bmatrix}, \quad v = \begin{bmatrix} 3/5 \\ 1/10 \\ 3/10 \end{bmatrix}. \quad (9)$$

After the third agent arrives, use the complete three-node graph and $c = 1/4$. Its averaging matrix is $P_3 = \frac{1}{4}I_3 + \frac{1}{4}\mathbf{1}_3\mathbf{1}_3^\top$. One update gives

$$P_3\mathcal{J}_{\beta,v}(W) = \begin{bmatrix} 7/15 & 5/24 & 13/40 \\ 1/3 & 41/120 & 13/40 \\ 7/15 & 13/60 & 19/60 \end{bmatrix}. \quad (10)$$

Remove the third agent proportionally. The retained row mass is $27/40$ for both original agents, and therefore

$$U = \mathcal{D}_{\{1,2\}}(P_3\mathcal{J}_{\beta,v}(W)) = \begin{bmatrix} 56/81 & 25/81 \\ 40/81 & 41/81 \end{bmatrix}. \quad (11)$$

All inputs, intermediate priorities and gains are admissible. The population and labels have returned to their initial values, but

$$m(W) = (1/2, 1/2), \quad m(U) = (16/27, 11/27). \quad (12)$$

By the mean criterion in Section II-A, subsequent ordinary consensus cannot restore the original limiting priority. This refutes unrestricted endpoint invariance even with only one intervening averaging step. It does not contradict the immediate identity (8).

B. A constructive characterization of a round trip

The final-column-sum criterion can be resolved into the inputs of a particular history. Fix $W \in \mathcal{S}_n^+$, incumbent insertion masses $\beta_i \in (0, 1)$, and a newcomer row

$$v^\top = (\gamma q^\top, 1 - \gamma), \quad 0 < \gamma < 1, \quad q \in \Delta_n^+.$$

Thus q is the newcomer's conditional priority on the original objectives. Let $T \in \mathbb{R}^{(n+1) \times (n+1)}$ be any entrywise nonnegative row-stochastic matrix: $T\mathbf{1}_{n+1} = \mathbf{1}_{n+1}$. It need not be connected, symmetric, doubly stochastic, or a power of a fixed consensus matrix. It mixes rows while the expanded population is fixed; the appended agent alone then departs. The characterization holds T and the retention masses fixed when varying q . Define

$$a = (1 - \beta_1, \dots, 1 - \beta_n, \gamma)^\top, \quad C = \begin{bmatrix} W \\ q^\top \end{bmatrix}, \quad (13)$$

$$d_i = \sum_{r=1}^{n+1} T_{ir} a_r, \quad \theta_r = \frac{a_r}{n} \sum_{i=1}^n \frac{T_{ir}}{d_i}, \quad 1 \leq r \leq n+1. \quad (14)$$

The coefficients d_i and θ_r depend only on the mixing and retention masses, before specifying q or forming the returned matrix.

Theorem 4.1 (Offline round-trip endpoint characterization). For the history above, all d_i are positive, $\theta_r \geq 0$, and $\sum_r \theta_r = 1$. The returned state $U = \mathcal{D}_{\{1, \dots, n\}}(T\mathcal{J}_{\beta,v}(W))$ belongs to $\mathcal{S}_n^+$ and satisfies

$$m(U) = \sum_{r=1}^n \theta_r w_r^\top + \theta_{n+1} q^\top. \quad (15)$$

Set $b^\top = m(W) - \sum_{r=1}^n \theta_r w_r^\top$. For fixed W, β, γ, T :

(i) If $\theta_{n+1} > 0$, the original endpoint is restored exactly when

$$q^\top = q_*^\top := \frac{b^\top}{\theta_{n+1}}. \quad (16)$$

Such a strictly positive newcomer priority exists if and only if every component of b is positive; when it exists it is unique. Its endpoint error for another admissible q is exactly

$$m(U) - m(W) = \theta_{n+1}(q - q_*)^\top. \quad (17)$$

(ii) If $\theta_{n+1} = 0$, the endpoint is restored if and only if $b = 0$, independently of q .

Proof. The first n columns of the joined matrix are $\text{diag}(a)C$, and its last column is $\mathbf{1}_{n+1} - a$. Each row of T is a nonnegative unit-sum row, so mixing preserves strict positivity and row stochasticity by convex combination. Since every $a_r \in (0, 1)$, its old-coordinate mass after mixing, $d_i = (Ta)_i$, lies in $(0, 1)$. Deleting the appended row and column and conditioning each survivor gives a positive probability row and hence $U \in \mathcal{S}_n^+$, with

$$U_{ij} = \frac{\sum_{r=1}^{n+1} T_{ir} a_r C_{rj}}{d_i}, \quad 1 \leq i, j \leq n.$$

Averaging these rows and interchanging the finite sums yields (15). The coefficients are nonnegative and

$$\sum_{r=1}^{n+1} \theta_r = \frac{1}{n} \sum_{i=1}^n \frac{\sum_r T_{ir} a_r}{d_i} = 1.$$

The endpoint equality is consequently $\theta_{n+1}q^\top = b^\top$. For positive θ_{n+1} this has the unique solution (16). Since the rows of W sum to one, $\sum_j b_j = 1 - \sum_{r=1}^n \theta_r = \theta_{n+1}$; hence the solution sums to one and belongs to Δ_n^+ exactly when b is entrywise positive. Subtraction gives (17). When the coefficient is zero, the equality reduces to $b = 0$. $\square$

This is an offline test using W, β, γ, T and the proposed q . Computing it generally requires global row information and the specified mixing. It does not implement an online distributed arrival protocol. A formal solution outside the strictly positive simplex is infeasible. With arbitrary T , zero newcomer influence need not preserve the incumbent mean; the condition $b = 0$ remains essential.

Consensus specialization. If the stay consists of ℓ updates of a fixed connected undirected consensus process, set $T = P^\ell$ using (4) in the expanded dimension. For $\ell \geq 1$, a newcomer edge and positive self-weights imply $\theta_{n+1} > 0$. For $\ell = 0$, $T = I$, $\theta = (1/n, \dots, 1/n, 0)$ and $U = W$. The general theorem does not require these stronger hypotheses.

Corollary 4.2 (Equal retention masses). In addition to the theorem hypotheses, suppose T is doubly stochastic. Let $\beta_i = \beta$ for every incumbent and let the newcomer assign its own objective the same mass $v_{n+1} = \beta$. Let $s = 1 - T_{n+1, n+1}$. If $s > 0$, the endpoint is restored if and only if

$$q^\top = \frac{\sum_{j=1}^n T_{n+1, j} w_j^\top}{s}. \quad (18)$$

This conditional priority is always strictly positive and feasible, for arbitrary $W \in \mathcal{S}_n^+$. For one-step Laplacian mixing $T = I - cL_G$ on a simple undirected graph, with $c > 0$ and a non-isolated newcomer, it is the arithmetic mean of the original priority rows of the newcomer's neighbors. If $s = 0$, every admissible q restores the endpoint under these equal-retention assumptions.

Proof. Now $a = (1 - \beta)\mathbf{1}_{n+1}$, so $d_i = 1 - \beta$. Double stochasticity of T gives $\theta_j = (1 - T_{n+1, j})/n$ and $\theta_{n+1} = s/n$. Substitution into (15) gives the exact error

$$m(U) - m(W) = \frac{1}{n} \left(s q^\top - \sum_{j=1}^n T_{n+1, j} w_j^\top \right).$$

If $s > 0$, its vanishing is (18). The weights in that expression are nonnegative and sum to one, so their combination of positive probability rows is positive and normalized. For $T = I - cL_G$, $T_{n+1, j} = c$ on each newcomer edge and zero otherwise, while $s = c \deg_G(n+1)$; cancellation yields the neighbor-row mean. If $s = 0$, the last row of T has zero entries in the first n columns, so the error vanishes for all q . $\square$

Corollary 4.2 permits disagreeing incumbent rows and can prescribe a newcomer priority different from their global mean. In the one-update case, those neighbors' pre-arrival priority rows give a locally interpretable formula, assuming coherent labels and access to those rows. This is not a complete membership protocol. For a longer prescribed stay, the full mixing matrix matters; unknown departure time and unmodeled graph changes are not covered. The earlier common-conditional-preference sufficient family is retained in Appendix D.

A restoring arrival whose effect depends on its stay. Take

$$W = \begin{bmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{bmatrix}, \quad P = \begin{bmatrix} 3/4 & 1/4 & 0 \\ 1/4 & 1/2 & 1/4 \\ 0 & 1/4 & 3/4 \end{bmatrix},$$

$$\beta_1 = \beta_2 = 1/3, \quad \gamma = 2/3. \quad (19)$$

For a one-update stay, the newcomer's only neighbor is agent 2, so choose $q^\top = w_2^\top = (1/4, 3/4)$, giving $v^\top = (1/6, 1/2, 1/3)$. Departure after one update returns

$$U = \begin{bmatrix} 5/8 & 3/8 \\ 3/8 & 5/8 \end{bmatrix}, \quad m(U) = m(W) = (1/2, 1/2).$$

The states differ, and neither incumbent agreement nor equality of q to the global mean was required. If the same newcomer instead stays for two updates, the returned mean is $(31/64, 33/64)$. For a two-update stay, (18) instead prescribes $q_*^\top = (1/3, 2/3)$, which restores $(1/2, 1/2)$. These identities are checked by exact arithmetic in the source package.

V. OPTIMIZATION INTERPRETATION AND LIMITS

On a fixed objective set, the final positive priority row μ defines $F_\mu(x) = \sum_j \mu_j f_j(x)$. Every attained global minimizer is Pareto efficient: a weak improvement of all objectives with one strict improvement would lower the positive weighted sum [20]. This interpretation does not establish convergence of any decision-update algorithm.

Changing the endpoint can change a unique optimizer. For $X = [0, 1]$, $f_1(x) = x^2$ and $f_2(x) = (x - 1)^2$, one has $F_\mu(x) = (x - \mu_2)^2 + \mu_2 - \mu_2^2$, so $x_\mu^* = \mu_2$. The counterexample in (12) selects $1/2$ or $11/27$ if the scalarized problem is solved exactly. Equal final populations alone do not fix that answer.

The characterization applies to the specified state, retention masses and mixing. Appendices C–D supply the standard finite-event continuation proof and secondary reset, composition and implementation details. Unknown-departure-time invariance and a distributed realization of the offline construction remain outside the claims.

<https://www.hortrame.com/swarm/>

The versioned package binds the supplied software; the live site is mutable. Optional synthetic link delays are handled by a simulated synchronization barrier, preserving the ordinary logical iteration. A controlled trial computes the newcomer priority for a prescribed consensus-power stay and checks the returned mean. These demonstrations do not establish asynchronous delay compensation or a decentralized membership protocol. Physical motion uses the precomputed epoch mean and is illustrative. This edition includes the three original screenshots.

No physical-robot validation of the membership-state formulas or current browser implementation is reported; that validation is planned separately. Implementation chronology is retained in the supporting provenance record.

Conditional fault-management interpretation. The admissible membership maps and finite-event continuation result do not require priority consensus before a transition. They can therefore supply state-transition semantics for a policy that deliberately treats sustained communication loss as departure

and later reconnection as re-entry. Re-entry requires a declared arrival state and coherent objective labels, not unqualified reuse of a stale vector. Fault detection, agreement on the active set, connectivity after removal, state transfer and the treatment of continuing failures remain separate obligations. The result does not guarantee decentralized fault tolerance or preservation of the original endpoint; restoration still requires the stated conditions. The browser's loss-free transport barrier does not implement such a fault-management policy.

VI. RESEARCH PROCESS AND VERIFICATION

The epistemological framework informs the ACGR methodology used to organize explicit claims, adversarial challenges, counterexamples, primary-source comparisons and reproducible checks. The companion records the evidence and reasons for revision [21]. This is a transparency record, not a methodological contribution of this paper. The full methodology remains under development and is intended for presentation in future research. The workflow does not replace proof or establish independent human acceptance.

APPENDIX A

FIXED-EPOCH PROOF AND STORAGE COORDINATES

The following supplies the standard proof of Lemma 2.1 under (4).

Proof. The off-diagonal entries of P are $cH_{ij} \geq 0$ and its diagonal entries are $1 - cd_i > 0$. Since $L\mathbf{1}_n = 0$ and $L = L^\top$, both row and column sums of P equal one. Each updated row is a convex combination of previous rows. The positive diagonal makes each positive previous entry contribute positively to the same updated row, proving positivity preservation.

For completeness, the required spectral bounds follow directly from

$$x^\top Lx = \sum_{\{i,j\} \in E} (x_i - x_j)^2 \geq 0.$$

Connectivity makes the kernel exactly $\text{span}\{\mathbf{1}_n\}$. Also

$$x^\top Lx \leq 2 \sum_{\{i,j\} \in E} (x_i^2 + x_j^2) = 2 \sum_i d_i x_i^2 \leq 2d_{\max} \|x\|_2^2.$$

Thus the eigenvalues satisfy $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n \leq 2d_{\max}$. The eigenvalues of P off the consensus subspace are $1 - c\lambda_j \in (-1, 1)$. Set $\rho = \max_{j \geq 2} |1 - c\lambda_j| < 1$. Orthogonal diagonalization yields the Euclidean contraction on $\mathbf{1}_n^\perp$; applying it column by column gives (5). Finally $\mathbf{1}_n^\top P = \mathbf{1}_n^\top$ proves mean preservation. $\square$

The proof is included to make the convergence premise explicit; it is a standard average-consensus argument, consistent with the established literature [19]. Strictness of the gain bound matters. On two connected nodes, $c = 1$ gives a swapping matrix and can produce periodic disagreement. The choice $c = 0$ also fails to converge from unequal rows.

A. Storage coordinates are not membership events

Let an ordered set V occupy n slots of an N -slot implementation. Write $S_V \in \{0, 1\}^{n \times N}$ for its selection matrix. Padding and extraction satisfy

$$S_V(S_V^\top W S_V)S_V^\top = W, \quad (20)$$

because $S_V S_V^\top = I_n$. This identity removes known inactive zero slots without changing the active state. It does not justify removing an active row or objective.

If Π consistently relabels agents and objectives, then $W' = \Pi W \Pi^\top$, $P' = \Pi P \Pi^\top$ and $P' W' = \Pi(PW)\Pi^\top$. Thus noncontiguous physical IDs are harmless when all arrays are indexed consistently. Stale IDs or inconsistent row/column ordering do not satisfy this identity.

APPENDIX B

SEQUENTIAL DEPARTURES AND RESET ORDER

Proposition B.1 (Consistent sequential departures). For nonempty ordered subsets $B \subseteq A \subseteq V$ and $W \in \mathcal{S}_n^+$, departure to A followed immediately by departure to B equals direct departure to B , with selections expressed in their current coordinates.

Proof. For $i, j \in B$, the two normalizations give

$$\frac{w_{ij} / \sum_{\ell \in A} w_{i\ell}}{\sum_{h \in B} w_{ih} / \sum_{\ell \in A} w_{i\ell}} = \frac{w_{ij}}{\sum_{h \in B} w_{ih}}.$$

Removed rows are the same in both constructions. This proves equality. No intervening averaging or priority event is allowed in this identity. $\square$

A. Priority resets and update order

For $v \in \Delta_n^+$ define

$$\mathcal{R}_{i,v}(W) = W + e_i(v - w_i)^\top. \quad (21)$$

This replaces row i with $v^\top$, so it preserves $\mathcal{S}_n^+$. So does a convex interpolation of that row with v using a coefficient in $[0, 1]$. Resetting before mixing, $P\mathcal{R}_{i,v}(W)$, and resetting after mixing, $\mathcal{R}_{i,v}(PW)$, both preserve the simplex but generally have different means.

An additive rule with the correction measured before mixing is a third operation:

$$\mathcal{H}_{i,v}(W) = PW + e_i(v - w_i)^\top. \quad (22)$$

Here and below this is an abstract ungated map; it is not a transcript of every archived command path.

Proposition B.2 (Exact admissibility and mean comparison). For $W \in \mathcal{S}_n$, $v \in \Delta_n$ and nonnegative doubly stochastic P , the rows of $\mathcal{H}_{i,v}(W)$ sum to one. This matrix is nonnegative if and only if

$$(PW)_{ij} + v_j - w_{ij} \geq 0 \quad \text{for every } j. \quad (23)$$

Its mean equals the mean of the pre-mixing reset:

$$m(\mathcal{H}_{i,v}(W)) = m(P\mathcal{R}_{i,v}(W)) = m(W) + \frac{1}{n}(v - w_i)^\top, \quad (24)$$

$$m(\mathcal{R}_{i,v}(PW)) = m(W) + \frac{1}{n}(v^\top - (PW)_{i,:}). \quad (25)$$

The pre- and post-mixing means agree exactly when $(PW)_{i,:} = w_i^\top$.

Proof. The correction has zero row sum. All other rows remain those of PW , hence are nonnegative. The altered row is nonnegative precisely under (23). Left multiplication by $n^{-1}\mathbf{1}_n^\top$, using $\mathbf{1}_n^\top P = \mathbf{1}_n^\top$, proves both mean formulas. Their difference is $n^{-1}((PW)_{i,:} - w_i^\top)$. $\square$

Time-varying-priority formulations motivate checking the domain of additive changes separately [2]. An exogenous declared-priority sequence is not identified here with the current-row correction in (22). The proposition is a general algebraic check; its statement and proof do not depend on a private implementation command.

APPENDIX C

A CONTINUATION COROLLARY FOR FINITE MEMBERSHIP CHANGES

The following statement records a consequence of Lemma 2.1. Once each event returns an admissible state, continuation follows by induction and convergence follows from the final fixed epoch. The particular departure and arrival formulas are not needed for that implication; their role is to supply explicit admissible semantics and determine the endpoint examined in Section IV.

Let $W_0 \in \mathcal{S}_{n_0}^+$. There are $R \geq 0$ event boundaries. Before event $r + 1$, epoch r executes $\ell_r \in \{0, 1, 2, \dots\}$ ordinary updates, where ℓ_r is finite. Define

$$W_{r+1} = E_{r+1}(P_r^{\ell_r} W_r), \quad r = 0, \dots, R-1. \quad (26)$$

The dimensions of P_r and W_r are $n_r \times n_r$, and E_{r+1} maps the current state to dimension n_{r+1} . All populations in this corollary are nonempty. Each P_r obeys (4), including the singleton convention. Every event is a finite, dimension-consistent composition of the operators in Section III, with admissible parameters; more generally it can be any explicitly defined map $\mathcal{S}_{n_r}^+ \rightarrow \mathcal{S}_{n_{r+1}}^+$.

After the last event, no further membership, priority or graph changes occur. Ordinary updates continue indefinitely with fixed P_R :

$$W_R(t) = P_R^t W_R, \quad t = 0, 1, 2, \dots \quad (27)$$

Event boundaries in (26) are atomic state transitions in the mathematical specification. Zero iterations between consecutive events are permitted. This formalizes “without waiting for consensus”; it does not impose a wall-clock deployment deadline.

Corollary C.1 (Admissible finite-event continuation). Under the preceding assumptions:

- (i) Every ordinary and post-event state is positive row-stochastic in its current dimension.
- (ii) No convergence or small-disagreement hypothesis is required before any event.
- (iii) In the final epoch,

$$W_R(t) \longrightarrow \mathbf{1}_{n_R} \mu_R, \quad \mu_R = m(W_R) \in (\Delta_{n_R}^+)^{\top}. \quad (28)$$

For $n_R \geq 2$, the error satisfies (5) with P_R, W_R and ρ_R . For $n_R = 1$ it is identically zero.

Proof. Lemma 2.1 preserves $\mathcal{S}_{n_r}^+$ through any finite number ℓ_r of ordinary updates. The event map then returns a state in $\mathcal{S}_{n_{r+1}}^+$. Induction over the finite list of events proves (i). Neither induction step refers to disagreement or convergence; this proves (ii). After event R , (27) is precisely a fixed-population process. Applying Lemma 2.1 gives (iii), including the stated bound. The mean of positive stochastic rows is itself positive stochastic. $\square$

The limit depends on the actual finite-time states supplied to the event operators. Those states can depend on all preceding graphs, gains, event parameters and epoch lengths. Eventual stationarity is substantive: finitely many membership changes alone would not imply (28) if externally forced priorities continued changing forever. Nor does the corollary supply a uniform convergence rate over arbitrary final graphs and gains. Even for a fixed graph, taking c toward zero can make ρ_R approach one.

A. A graph bound that remains valid after changes

If at most $N \geq 2$ agents can be active and graphs are simple, then $d_{\max} \leq n - 1 \leq N - 1$. Any fixed gain

$$0 < c < \frac{1}{N-1} \quad (29)$$

therefore satisfies the corollary’s gain condition in every nonsingleton epoch. Alternatively, recompute c from each committed graph. The graph must also be connected; the gain bound alone does not provide connectivity. The bound is a sufficient design condition, not a claim of fastest convergence or necessity.

B. Empty populations and repeated changes

There is no probability distribution on zero objectives in this formulation. A departure of the last agent leaves the corollary’s state space. A system can specify an inactive mode and initialize a fresh state when agents return; that is a restart with new data, not a continuation preserving an empty state’s mean. A two-to-one departure is covered and yields [1]. One-to-two arrival is covered by (7).

If events continue indefinitely, (28) makes no assertion. Tracking and open-system stability then require appropriate definitions and additional assumptions, as in the cited open-system literature [9], [11], [12].

APPENDIX D

ADDITIONAL ENDPOINT AND SCALARIZATION DETAILS

A. Exact event means

Suppose no further events follow a departure to A , with $m = |A|$. For each surviving objective j ,

$$\mu_{A,j} = \frac{1}{m} \sum_{i \in A} \frac{w_{ij}}{\sum_{h \in A} w_{ih}}. \quad (30)$$

This is an average of conditional rows. Generally it is not obtained by conditioning the pre-event population mean: both the set of averaging agents and the row-dependent denominators change.

After a single arrival with no further events, the limiting priorities are

$$\mu_{J,j} = \frac{1}{n+1} \left(\sum_{i=1}^n (1 - \beta_i) w_{ij} + v_j \right), \quad 1 \leq j \leq n, \quad (31)$$

$$\mu_{J,n+1} = \frac{1}{n+1} \left(\sum_{i=1}^n \beta_i + v_{n+1} \right). \quad (32)$$

These formulas follow by taking the row mean of (6) and (7) and then applying Corollary C.1. They specify the target rather than assuming that it is unaffected.

Proposition D.1 (When two histories have the same limit). Consider two histories satisfying Corollary C.1, ending on the same ordered label set with final initial states U and V of order m . Their limiting priority matrices are equal if and only if

$$\mathbf{1}_m^\top U = \mathbf{1}_m^\top V. \quad (33)$$

Their final connected undirected graphs and admissible gains need not be identical.

Proof. The two limits are $\mathbf{1}_m m(U)$ and $\mathbf{1}_m m(V)$. They are equal exactly when their common rows agree, which is (33) after multiplication by m . $\square$

In particular, when a history returns to the original ordered population, its limiting priorities equal those of the event-free run if and only if the returned state's row mean equals the original row mean. Intermediate event means need not be equal. The immediate identity in (8) already demonstrates this distinction: insertion may change the dimension and mean before deletion restores the original state. Simplex preservation alone does not imply endpoint preservation. When final label sets differ, one must first define a comparison between different objective spaces; literal equality of the vectors is not meaningful.

B. A sufficient special case for delayed removal

The negative result does not exclude useful conditional identities.

Proposition D.2 (Common conditional preferences). Fix $u \in \Delta_n^+$. Suppose every row of an $(n+1)$ -agent state has the form

$$z_i^\top = (a_i u^\top, 1 - a_i), \quad 0 < a_i < 1. \quad (34)$$

Any finite sequence of left multiplications by nonnegative row-stochastic matrices of order $n+1$ preserves this form. Proportional removal of the appended agent afterward yields $\mathbf{1}_n u^\top$ for the survivors.

Proof. For a row-stochastic matrix B , updated row i has old-coordinate portion $(\sum_j B_{ij} a_j) u^\top$ and last coordinate $1 - \sum_j B_{ij} a_j$. A convex combination of finitely many numbers in $(0, 1)$ remains in $(0, 1)$. Induction proves preservation through finitely many such updates. Departure divides the old-coordinate portion by its row sum, recovering $u^\top$ in each surviving row. $\square$

Thus arrival into an already agreed state can preserve that agreement after later removal if the newcomer's conditional preferences on the original objectives agree as well. This is a stated sufficient condition, not a justification for arbitrary arrivals with arbitrary preferences.

C. What the priorities imply for optimization

Let the final agents own real-valued objectives f_j on a common nonempty feasible set X . The final priority row μ defines

$$F_\mu(x) = \sum_{j=1}^m \mu_j f_j(x). \quad (35)$$

If the f_j are continuous and X is compact, F_μ attains a minimum. Positive weights give the following elementary scalarization fact, also standard in convex-optimization treatments [20].

Proposition D.3 (Scalarization interpretation). Every global minimizer of (35), for strictly positive μ , is Pareto efficient on X .

Proof. If $y \in X$ weakly improves every objective over a minimizer x and strictly improves at least one, positivity implies $F_\mu(y) < F_\mu(x)$, a contradiction. Convexity is not needed for this implication. $\square$

Changing μ can change even a unique optimizer on an unchanged objective set. For $X = [0, 1]$, $f_1(x) = x^2$ and $f_2(x) = (x - 1)^2$, completing the square gives

$$F_\mu(x) = (x - \mu_2)^2 + \mu_2 - \mu_2^2, \quad x_\mu^* = \mu_2. \quad (36)$$

The two histories in (12) therefore select $1/2$ and $11/27$, respectively, if that scalarized problem is solved exactly. This is an objective-level counterexample, not a simulated robot trajectory.

Corollary C.1 proves no convergence of decision iterates to these minimizers. Such a conclusion needs an explicit decision update, its state and communication assumptions, and a valid convergence proof for the final objective. Restarting a theorem is legitimate only after all its hypotheses and the entire algorithmic state are matched. Here no full-MOO convergence theorem is imported, and no physical tracking, collision avoidance or rendezvous guarantee is deduced from agreement of W .

APPENDIX E

IMPLEMENTATION SEMANTICS AND SCOPE

A. Centralized rows and local cached rows

The motivating supplied MATLAB simulation stored the full priority matrix centrally and updated its rows from one shared old state. The subsequent modular QBP implementation in Simulink/QUARC ran both in simulation and on physical robots from the Quanser QBot Platform family before Naderi's delay refinement; it used persistent per-agent state, an active-label ordering and cached neighbor rows. Membership-transition handling was unfinished in the saved implementation. The supplementary provenance records these source-level facts, their precise artifact bindings and their limits; they are not assumptions in a mathematical proof.

The algebra relevant to the specification is simple. A local row update of the form

$$w_i^+ = (1 - c d_i) w_i + c \sum_{j \in N_i} \hat{w}_{ij} \quad (37)$$

is row i of (4) when every cache $\hat{w}_{ij}$ contains row j of the same previous globally indexed state. A local star contains the

incident edges needed for that row; edges between its neighbors are unnecessary for this calculation. Asynchronous caches or inconsistent label sets do not satisfy that identity merely because the same formula is written locally. No preserved source version is certified here as a realization of the complete event specification.

B. An implementation contract for the proved semantics

An implementation claiming to realize (26) must supply a coherent active-label ordering and a well-defined event state. In particular:

- 1) Commit the event and current state under an agreed membership interpretation; incompatible local views lie outside the abstract transition.
- 2) Select a nonempty survivor set, verify each retained row's positive surviving mass, and normalize once. An empty population requires a separate inactive mode.
- 3) Insert a fresh objective coordinate with declared $0 < \beta_i < 1$ and $v \in \Delta_{n+1}^+$. Specify coordinate semantics if physical identifiers are reused.
- 4) Order priority resets explicitly before or after mixing; an unchecked additive correction is a different operator.
- 5) Commit a connected simple undirected graph and an admissible positive gain, or $P = [1]$ for a singleton. Resume updates only with matching row and column labels.

This is a state contract, not a newly implemented membership protocol. Atomicity may require coordination; Corollary C.1 does not establish instantaneous handling of an unannounced physical failure. Acknowledgement-based open consensus explicitly addresses related accounting and protocol obligations [13], [14].

C. Representative browser parameters

The online application exposes a common incumbent insertion mass, defaulting to $\beta_i = 1/(n+1)$, and samples a positive newcomer row. Its self-priority reset shares the remaining mass uniformly. These are representative parameterizations, not all admissible maps; uniform insertion mass does not mean endpoint neutrality. In particular, its ordinary manual arrival does not enforce the equal-retention hypothesis and prescribed conditional row of Corollary 4.2.

The epoch counter advances only for membership and priority events. Graph/gain changes preserve W and can be modeled as identity-state boundaries; a final fixed epoch is still required.

D. Synthetic transport and a controlled restoration experiment

The browser optionally schedules one old-row message on each directed communication link, with uniform or independently varied finite delivery times. Each simulated logical round waits until all required messages from its common old state have arrived before committing (4). The barrier therefore changes elapsed simulation time, not the state sequence indexed by committed rounds. Membership, priority and graph/gain

changes discard pending obsolete rounds. This is a loss-free, centrally coordinated simulation with coherent event boundaries; it is not Naderi's asynchronous delay-correction algorithm or a general solution for delayed open membership. Latency parameters use simulation time, which freezes when paused. Message loss, decentralized agreement on epochs and physical transport are outside this demonstration.

A separate controlled arrival-departure trial instantiates Corollary 4.2. It saves the incumbent state, chooses a common insertion/self mass, forms the expanded graph and fixes a stay of ℓ committed updates. It computes the conditional newcomer row from the last row of P^ℓ , runs exactly ℓ updates, then removes that newcomer and compares the returned mean with the saved mean. Manual membership, priority or graph/gain changes interrupt the trial. Delay and motion settings may vary without changing the prescribed logical mixing. Numerical equality is reported with a displayed error and tolerance, not as a new proof. The normal manual-arrival controls still sample a positive newcomer row. The controlled trial covers equal retention and consensus-power mixing, not every input of the general theorem or an unknown departure time.

APPENDIX F

EXACT EXAMPLES AND REPRODUCIBLE ILLUSTRATION

The following algebraic examples use the specified synchronous maps, not the robot system. The accompanying Python program checks them in rational arithmetic and generates the table and curve; floating-point conversion is limited to the displayed norm and plot. Appendix F-C displays the states of the restoring round trip. The separate browser visualization remains available online and in supplementary images.

A. Reset order and loss of nonnegativity

Take

$$P = \begin{bmatrix} 3/5 & 2/5 \\ 2/5 & 3/5 \end{bmatrix}, \quad W = \begin{bmatrix} 9/10 & 1/10 \\ 1/10 & 9/10 \end{bmatrix},$$

$$v = \begin{bmatrix} 1/10 \\ 9/10 \end{bmatrix}, \quad i = 1. \quad (38)$$

Then $(PW)_{1,:} = (29/50, 21/50)$ and $v^\top - w_1^\top = (-4/5, 4/5)$. The additive map gives

$$\mathcal{H}_{1,v}(W) = \begin{bmatrix} -11/50 & 61/50 \\ 21/50 & 29/50 \end{bmatrix}. \quad (39)$$

It has unit row sums but is not a priority matrix. Meanwhile

$$\begin{aligned} m(PR_{1,v}(W)) &= (1/10, 9/10), \\ m(\mathcal{R}_{1,v}(PW)) &= (13/50, 37/50). \end{aligned} \quad (40)$$

Both reset compositions preserve the simplex; their limiting priorities differ. The additive map has the first mean but an inadmissible intermediate state. No claim of equivalence between all three operations is made.

TABLE I
EXACT MEANS AT THE ILLUSTRATIVE EVENT BOUNDARIES.

Transition	State	Mean priority
0	Initial	(1/2, 1/2)
3	Arrival	(19/45, 23/90, 29/90)
7	Departure	(808/1301, 493/1301)
10	Reset	(10677/26020, 15343/26020)
40	Final sample	(10677/26020, 15343/26020)




Fig. 1. Illustrative finite-event process. Arrival occurs at transition 3, departure at 7 and row reset at 10. After transition 10, each ordinary step halves the disagreement; the figure illustrates the proved final-epoch contraction, not delay robustness or historical robot performance.

B. A finite sequence with changes before consensus

Initialize the two-agent matrix in (9). Perform two ordinary complete-graph updates with $c = 1/4$, then insert a third agent using the same β and v as in (9). Perform three complete-three-agent updates with $c = 1/4$, remove the third agent proportionally, and perform two complete-two-agent updates with $c = 1/4$. Reset the first row to $(1/5, 4/5)$ at that event boundary, then continue ordinary two-agent consensus for 30 iterations. Each event thus has an explicit order; no event waits for exact consensus.

The abscissa in Fig. 1 counts state transitions, including events as separate transitions. Its error is measured against the mean of the current epoch, not a fictitious fixed mean across changing dimensions. Events can change both the state and that reference mean. The final mean and selected checkpoint values are generated directly from the exact recurrence in Table I.

Exact checks also cover immediate insertion/deletion, sequential departures, relabeling, singleton states and undefined inputs. General claims rely on the printed proofs.

C. A round trip with a restored endpoint

Figure 2 follows the exact data in (19): an arrival, one expanded-population update, and departure of that same agent. It displays the mathematical objects affected by membership. In each panel μ is the state's own row mean and $\varepsilon = \|W - 1m(W)\|_F$ in that panel's dimension.

The browser simulation [18] provides interactive arrivals, departures and resets. Its motion target is computed directly from the epoch mean, independently of the local rows in evolution. Three original beginning/middle/celebration screenshots are preserved in Appendix I and are not used as membership or physical-convergence evidence.

<https://www.hortrame.com/swarm/>

APPENDIX G REPRODUCIBILITY AND VERSIONED SOFTWARE

Upload the source package to Overleaf and compile `main.tex` with pdfLaTeX. Pagination can vary with package versions. No shell escape is required. All mathematical assumptions and proofs appear in the note. Detailed implementation chronology, the author account and private artifact identities are supplied separately under `supplementary/` and `evidence/`; no result depends on access to those archives.

To regenerate the numerical evidence, run `python reproduce.py` from the project directory using Python 3.9 or newer. It regenerates the original trace and examples and runs two standard-library modules. `roundtrip_checks.py` compares direct join-mix-depart execution with the end-point formulas in 1,152 rational consensus-power cases. `general_mixing_checks.py` adds 768 cases with arbitrary row-stochastic mixing, including nonsymmetric, reducible, zero-diagonal and permutation matrices. The added cases check zero newcomer influence both with and without restoration, and exhibit failure of the equal-retention shortcut without double stochasticity. The 1,920 cases are finite cross-checks, not an exhaustive parameter-domain proof. Generated files are included, so Overleaf needs no Python.

The numbered author bundle includes an offline simulation package, a website package and the exact full PDF. The website package stores its public files under `swarm/` to preserve the company homepage; the offline entry point is `simulation/index.html`. Both include the bilingual interface and the English manuscript. The English and French interfaces execute the same priority and motion algorithms. The bundle's versioned SHA-256 sidecar identifies the delivered components, including source, PDFs and component ZIPs. The validation ledger records exact runtime hashes, dates and whether each check was newly run or inherited.

This binds the reviewed bytes; it is not a public DOI deposit. The live HORTAME URL is mutable and may later serve different code. A permanent public software deposit remains a publication task; no repository identifier or license is invented here.

APPENDIX H SCOPE OF VERIFICATION

Analytic proofs establish the stated claims. Exact rational tests protect printed examples and compare two computations of the round-trip endpoint. Software checks concern the stated priority implementation and finite controller scenarios; physical images are not evidence of distributed discovery of a motion target. The earlier browser-test report is dated separately from later non-browser and packaging checks. The authoring pass does not constitute proof-assistant certification, independent human review or publication acceptance. New results do not by themselves settle novelty or journal suitability.

APPENDIX I SUPPLEMENTARY VISUAL DEMONSTRATION

Figure 3 shows three stages of one fixed twelve-agent run of the public simulation on September 7, 2026, with no




Fig. 2. The population and original mean return, although the priority state changes and the intermediate mean is in a different objective space. Subsequent admissible fixed-population consensus converges to rows $(1/2, 1/2)$. This is the positive case of Corollary 4.2, with exact arithmetic rather than a physical-controller trace.

membership or priority event. The images illustrate the separate physical controller; the main membership figure is Fig. 2. Here k counts priority updates and ε is the displayed Frobenius disagreement. Capture identities and original data are supplied in `evidence/simulation_capture.json`.

ACKNOWLEDGMENT

AI agents assisted drafting and refinement throughout Sections I–VI, Appendices A–I, the bibliography and this acknowledgment, including mathematical formulation, proof construction and adversarial examination. It also assisted development of the accompanying verification code and simulation implementation. It is not an author. Simulation material in this manuscript is explicitly illustrative and is not used as proof of the mathematical claims. Human final review and acceptance remain pending; the human author is responsible for the manuscript. The previously displayed future acceptance paragraph is preserved in the revision archive and is not asserted as a completed event in this preprint.

REFERENCES

- [1] M. J. Blondin and M. Hale, “A decentralized multi-objective optimization algorithm,” *Journal of Optimization Theory and Applications*, vol. 189, no. 2, pp. 458–485, 2021. [Online]. Available: <https://doi.org/10.1007/s10957-021-01840-z>
- [2] S. Naderi and M. Blondin, “A distributed multi-objective optimization algorithm with time-varying priorities for multi-agent systems,” *Research Square*, version 1, Aug. 2025, posted Aug. 26, 2025. [Online]. Available: <https://doi.org/10.21203/rs.3.rs-7292808/v1>
- [3] S. Naderi, M. J. Blondin, and S. Roy, “Consensus in multi-agent systems with uniform and nonuniform communication delays,” arXiv:2603.16523v1, Mar. 2026. [Online]. Available: <https://arxiv.org/abs/2603.16523v1>
- [4] G. Oliva, M. Franceschelli, A. Gasparri, and A. Scala, “A sum-of-states preservation framework for open multiagent systems with nonlinear heterogeneous coupling,” *IEEE Transactions on Automatic Control*, vol. 69, no. 3, pp. 1991–1998, 2024. [Online]. Available: <https://doi.org/10.1109/TAC.2023.3317728>
- [5] B. Zhou, J. H. Park, Y. Yang, Y. Jiao, and R. Hao, “Dynamic weighted average consensus of open time-varying multi-agent systems on time scales via sampled-data impulsive communication,” *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 5, pp. 4330–4343, 2024. [Online]. Available: <https://doi.org/10.1109/TNSE.2024.3417386>
- [6] S. Jiang, X. Wu, D. Li, Z. Zhang, and Y. Zhang, “Sum-of-values conservation on strongly connected open systems,” *Communications in Nonlinear Science and Numerical Simulation*, vol. 162, 2026, Art. no. 110404. [Online]. Available: <https://doi.org/10.1016/j.cnsns.2026.110404>
- [7] J. M. Hendrickx and S. Martin, “Open multi-agent systems: Gossiping with deterministic arrivals and departures,” in *Proc. 54th Annual Allerton Conference on Communication, Control, and Computing*, 2016, pp. 1094–1101. [Online]. Available: https://perso.uclouvain.be/julien.hendrickx/availablepublications/Hendrickx_Martin_OpenMA_Allerton.pdf
- [8] R. Vizuete, C. Monnoyer de Galland, P. Frasca, E. Panteley, and J. M. Hendrickx, “Trends and questions in open multi-agent systems,” in *Hybrid and Networked Dynamical Systems*, R. Postoyan, P. Frasca, E. Panteley, and L. Zaccarian, Eds. Cham, Switzerland: Springer, 2024, vol. 493, pp. 219–252. [Online]. Available: https://doi.org/10.1007/978-3-031-49555-7_10
- [9] M. Franceschelli and P. Frasca, “Stability of open multiagent systems and applications to dynamic consensus,” *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 2326–2331, 2021. [Online]. Available: <https://doi.org/10.1109/TAC.2020.3009364>
- [10] R. Vizuete and J. M. Hendrickx, “Consensus on open multi-agent systems over graphs sampled from graphons,” in *Proc. 2025 IEEE 64th Conference on Decision and Control*, 2025, pp. 5282–5289. [Online]. Available: <https://doi.org/10.1109/CDC57313.2025.11312464>
- [11] J. M. Hendrickx and M. G. Rabbat, “Stability of decentralized gradient descent in open multi-agent systems,” arXiv:2009.05445v1, Sep. 2020. [Online]. Available: <https://arxiv.org/abs/2009.05445v1>
- [12] D. Deplano, N. Bastianello, M. Franceschelli, and K. H. Johansson, “Optimization and learning in open multiagent systems,” *IEEE Transactions on Automatic Control*, vol. 71, no. 6, pp. 3864–3879, 2026. [Online]. Available: <https://doi.org/10.1109/TAC.2026.3653296>
- [13] C. N. Hadjicostis and A. D. Domínguez-García, “Distributed average consensus in open multi-agent systems,” in *Proc. 2024 IEEE 63rd Conference on Decision and Control*, 2024, pp. 3037–3042. [Online]. Available: <https://doi.org/10.1109/CDC56724.2024.10886031>
- [14] E. Makridis, A. Grammenos, G. Oliva, E. Kalyvianaki, C. N. Hadjicostis, and T. Charalambous, “Average consensus over directed networks in open multi-agent systems with acknowledgement feedback,” arXiv:2409.08634v1, Sep. 2024. [Online]. Available: <https://arxiv.org/abs/2409.08634v1>
- [15] C. Wang, X. Wu, Y. Zhang, J. Lin, S. Zhang, and D. Li, “Consensus on




Fig. 3. Three genuine captures from the same public-application run. The motion controller uses the computed epoch mean; trails do not show distributed discovery of its target. Celebration requires both numerical priority agreement and physical settling. This fixed-population illustration is not evidence for a membership or physical-convergence theorem.

- open networked multi-agent systems,” *Systems & Control Letters*, vol. 214, Jul. 2026, Art. no. 106470. [Online]. Available: <https://doi.org/10.1016/j.sysconle.2026.106470>
- [16] J. J. Egozcue and V. Pawlowsky-Glahn, “Subcompositional coherence and a novel proportionality index of parts,” *SORT*, vol. 47, no. 2, pp. 229–244, 2023. [Online]. Available: <https://doi.org/10.57645/20.8080.02.7>
- [17] 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, Mar. 2026. [Online]. Available: <https://arxiv.org/abs/2603.08038v1>
- [18] G. Bolduc, “Rendezvous: Swarm lab,” HORTAME, interactive browser simulation, 2026, accessed Sep. 7, 2026. [Online]. Available: <https://www.hortame.com/swarm/>
- [19] R. Olfati-Saber, J. A. Fax, and R. M. Murray, “Consensus and cooperation in networked multi-agent systems,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007. [Online]. Available: <https://doi.org/10.1109/JPROC.2006.887293>
- [20] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge University Press, 2004. [Online]. Available: <https://web.stanford.edu/~boyd/cvxbook/>
- [21] G. Bolduc, “Epistemological method and revision history for the priority-state manuscript,” HORTAME, online research companion, 2026, work in progress; accessed Sep. 7, 2026. [Online]. Available: <https://www.hortame.com/swarm/epistemology/>
- [22] HORTAME Systems, “AI-assisted research tooling,” 2026, software.