

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

Guillaume Bolduc

Preprint (currently being written) — September 7, 2026

Version anglaise. Pour obtenir la version française lorsqu’elle sera disponible, sélectionnez « fr » sur la page du préprint, puis téléchargez le PDF français. La traduction française n’est pas encore disponible.

Working preprint. Not ready for formal review; an informal early look is welcome.

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.

Keywords: open multi-agent systems; priority-state dynamics; dynamic membership; stochastic matrices; history-dependent consensus endpoints.

1 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, Sec. II-B, p. 4332]. 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 [5, Def. 3 and Sec. III-A]; 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 [5, Thm. 1 and Rem. 8]; 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, Sec. II]; 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, Secs. 5–7]. 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.

2 Priority state and ordinary evolution

For $n \geq 1$, define

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

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

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.

2.1 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 the accompanying supporting material, 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.

3 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.

3.1 Departure

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

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

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.

3.2 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.

4 Endpoint restoration after a round trip

4.1 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 2, 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).

4.2 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 $s = c \deg_G(n+1)$, where $\deg_G(n+1)$ is the newcomer's graph degree; 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 the supporting material.

A restoring arrival whose effect depends on its stay. Take

$$W = \begin{bmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{bmatrix}, \quad P = \begin{bmatrix} 3/4 & 1/4 & 0 \\ 1/4 & 1/2 & 1/4 \\ 0 & 1/4 & 3/4 \end{bmatrix}, \quad \beta_1 = \beta_2 = 1/3, \quad \gamma = 2/3. \quad (19)$$

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.

5 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, Sec. 4.7]. 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. The accompanying supporting material supplies 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.

[Open interactive simulation ↗](https://www.hortrame.com/swarm/)
`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. The full 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.

6 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.

Acknowledgments and AI disclosure

OpenAI Codex [22] assisted drafting and revision of the abstract, Sections 1–6 and supporting material, including mathematical formulation and proof text; literature and archive analysis; generation of reproduction/test code and figures; and development of the browser implementation and bilingual companion. The three simulation panels are recorded browser captures, not generated raster illustrations. AI-assisted adversarial review informed revisions; received reports remain separately identified in the companion. In the original internship implementation, AI assistance concerned coding, not research direction or implementation decisions. The earlier supplied algorithm and implementation records were inputs to the later research process. Responsibility for the manuscript and its claims remains with the author; independent human mathematical approval is not asserted. Formal disclosure must be checked against the selected venue’s requirements at submission.

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 August 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, preprint, version 1, March 17, 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; assigned to the November 2026 issue. [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*, ser. Lecture Notes in Control and Information Sciences, 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 (CDC)*, 2025, pp. 5282–5289, final proceedings record. [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, author preprint. [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, author preprint arXiv:2501.16847v1 first posted January 28, 2025. [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 (CDC)*, 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, preprint, first posted September 13, 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 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. Pawłowsky-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, preprint, version 1, March 9, 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 September 7, 2026. [Online]. Available: <https://www.hortame.com/swarm/epistemology/>
- [22] OpenAI, “Codex,” 2026, software; accessed September 7, 2026. [Online]. Available: <https://openai.com/codex/>