Neuromorphic Decision Dynamics and Control for Provably Safe and Deadlock-Free Multi-Agent Navigation in Congested Environments

Himani Sinhmar[1*], Vaibhav Srivastava[2], and Naomi Ehrich Leonard[1]

***Abstract*— In congested multi-agent navigation, decisions and physical actions are tightly coupled because each agent's motion changes the evidence and options available to others. Safe deconfliction requires this coupled system to avoid collisions while establishing a traversal order that prevents deadlock and livelock. We present a constructive design principle for making this collective behavior a tunable, controllable, and certifiable emergent property of local physical interaction. Our Neuromorphic Decision Dynamics (NDD) controller is fully decentralized, path-preserving, and event-based. Each agent converts anonymous egocentric relative-motion observations into continuous-time conflict evidence; continuous-time decisions modulate traversal speed, and the resulting motion cues update neighboring decisions. This recurrent decision–action–perception loop forms a lightweight, fast interaction layer beneath a slower planner. When conflict becomes safety-critical, purely local feedback tunes the onset of an emergent group-level bifurcation in which collective indecision loses stability, small asymmetries amplify into differentiated motion, and temporal separation and an acyclic traversal order emerge without communication or centralized coordination. As conflicts clear, the same feedback smoothly restores nominal motion. We derive sensing, feedback, tracking, and actuation conditions that guarantee continuous safety, finite-time order formation and persistence, deadlock- and livelock-free progress, and robustness to bounded errors and time-varying interaction topologies, independently of population size. In randomized simulations with 4–32 Dubins agents, NDD achieved 100% safe-agent success and outperformed ORCA, LivePoint, and decentralized CBF-QP in safe throughput. It also completed a dense ten-agent near-antipodal exchange while preserving every path and maintaining the required clearance.**

I. Introduction

Multi-robot systems increasingly undertake complex real-world tasks in crowded spaces, where decisions and physical actions are tightly coupled: each agent's motion changes its neighbors' observations, options, and responses. This coupling intensifies at narrow corridors, intersections, merges, and other bottlenecks, where pedestrians, vehicles, and robots compete for limited space without a common protocol.

Agents engaged in these encounters must coordinate to avoid collisions and sustain progress. Deadlock appears as mutual yielding [1] or motion freezing when every forward action seems unsafe [2]; livelock appears as repeated adjustment without progress [3]. These failures arise when agents do not establish a consistent *traversal order* through contested space. Even sequences of locally feasible, safe actions can lead to deadlock or livelock [4]. We call collective behavior that preserves both safety and liveness *safe deconfliction*.

To our knowledge, establishing guarantees for decentralized safe deconfliction without preassigned traversal priorities remains an open challenge, even when communication is available. In decentralized settings, safe deconfliction requires each agent to account continuously for the evolving interaction structure: which neighbors are in conflict, how strongly each conflict should shape its response, and which agent should enter each contested region first. We provide a constructive solution with formal guarantees using decentralized nonlinear decision dynamics that exploit this structure online. The guarantees apply to arbitrarily large, time-varying populations with no inter-agent communication, partial observations of neighboring actions, and bounded actuation.

Many prediction-based coordination methods separate interaction assessment from movement decisions: an agent first predicts its neighbors' intended behavior and then selects an action conditioned on those predictions [2], [5]. This predict-then-act architecture reduces the coupled interaction to a single-agent control problem and assumes that the interaction structure can be determined before action and remain valid as agents respond. Existing methods also frequently separate safety from liveness. Reciprocal collision avoidance and control barrier functions enforce locally admissible motion [6], [7]; specialized constraints, detection schemes, and recovery policies support liveness [8], [9]; scheduling and priority methods share precedence information [10]; and dynamic games and coupled optimization model neighboring objectives and responses [11]. Dense encounters challenge these decompositions because each selected action reshapes the interaction structure. A neighbor's response may also reflect conflicts outside an agent's field of view. Independently formed traversal orders can disagree, conservative forecasts can eliminate all forward motion, and near-symmetric encounters can leave several orders equally plausible. Enumerating joint responses becomes combinatorially difficult as the group grows [12], [13].

Ongoing interaction generates transient, endogenous *coordination cues* through small speed changes and differences in projected arrival time. Every response changes the evidence observed by nearby agents, tightly coupling decision and motion. This motivates our scientific question: *Can a fully decentralized feedback controller exploit these cues to generate provably robust safe deconfliction in interacting groups of arbitrary and changing size?* Our thesis is that

[1]Himani Sinhmar and Naomi Ehrich Leonard are with the Department of Mechanical and Aerospace Engineering, Princeton University, NJ, USA.
[2]Vaibhav Srivastava is with the Department of Electrical and Computer Engineering, Michigan State University, MI, USA.
[*]Corresponding author: himani.sinhmar@princeton.edu.

safe deconfliction must emerge from closed-loop coupling among decisions, physical actions, and transient motion cues. By making deadlock and livelock unstable by construction, the controller enables agent motion to drive coordinated departure without explicit detection or repair. Jankovic *et al.* identified instability of unresolved equilibria as a mechanism for deadlock avoidance in small-group interactions and posed the constructive design of this mechanism under decentralized local feedback for arbitrary groups as an open problem [14], [15].

We address this open problem by introducing *Neuromorphic Decision Dynamics* (NDD), a continuous-time system that couples decision formation and path-preserving speed control to create controlled instability across an interacting group. NDD makes safe deconfliction a *controllable emergent property* of local physical interaction. Each agent operates independently using anonymous, egocentric measurements of nearby relative positions and velocities. NDD requires no communication, global information, common clock, synchronized updates, or models of neighboring agents. Within a multi-time-scale autonomy architecture, NDD acts as a fast interaction layer: a slower planner supplies each path, and NDD modulates only traversal speed as conflict evidence arises and subsides.

For arbitrary, time-varying groups, NDD creates a traversal order and temporal separation through a controlled group-level instability. The controller sets an evidence threshold above which unresolved states—indecision, deadlock, and livelock—lose stability. Above this threshold, arbitrarily small, interaction-relevant asymmetries select a stable *differentiated* motion pattern (Fig. 1(c)). This excursion resembles the response of an excitable neural dynamical system: decision states rapidly depart from zero and remain differentiated while conflict persists. Earlier-arriving agents accelerate and later-arriving agents slow, increasing temporal separation. The growing separation weakens conflict evidence, returns decision states to rest, and restores nominal motion. This self-limiting feedback adapts as evidence and group membership change.

We derive explicit conditions on sensing range, decision and control gains, and actuation limits that guarantee continuous safety, finite-time acyclic order formation, deadlock- and livelock-free progress, and recovery of nominal motion under bounded sensing and actuation errors. We demonstrate scalability in randomized and dense congestion simulations and benchmark NDD against ORCA, LivePoint, and decentralized CBF-QP [6], [7], [9].

II. Decentralized Safe Deconfliction

A. NDD and decentralized control

NDD augments each agent's nominal path follower with a decision state z_i and an adaptive attention state u_i. Agent i moves along a smooth prescribed path with position $\boldsymbol{q}_i$, speed v_i, and nominal speed v_i^{nom}. Its anonymous egocentric observation $\mathcal{X}_i$ contains its own motion and the relative motion of sensed neighbors. Each agent continuously constructs local conflict evidence $h_i(\mathcal{X}_i)$, which drives the controller:

$$\tau_z \dot{z}_i = -\rho_z z_i + S(u_i h_i(\mathcal{X}_i)), \tag{1a}$$

$$\tau_u \dot{u}_i = -\rho_u u_i + u_0 + \kappa_u z_i^2, \tag{1b}$$

$$v_i^{\mathrm{cmd}} = v_i^{\mathrm{nom}} + V(z_i). \tag{1c}$$

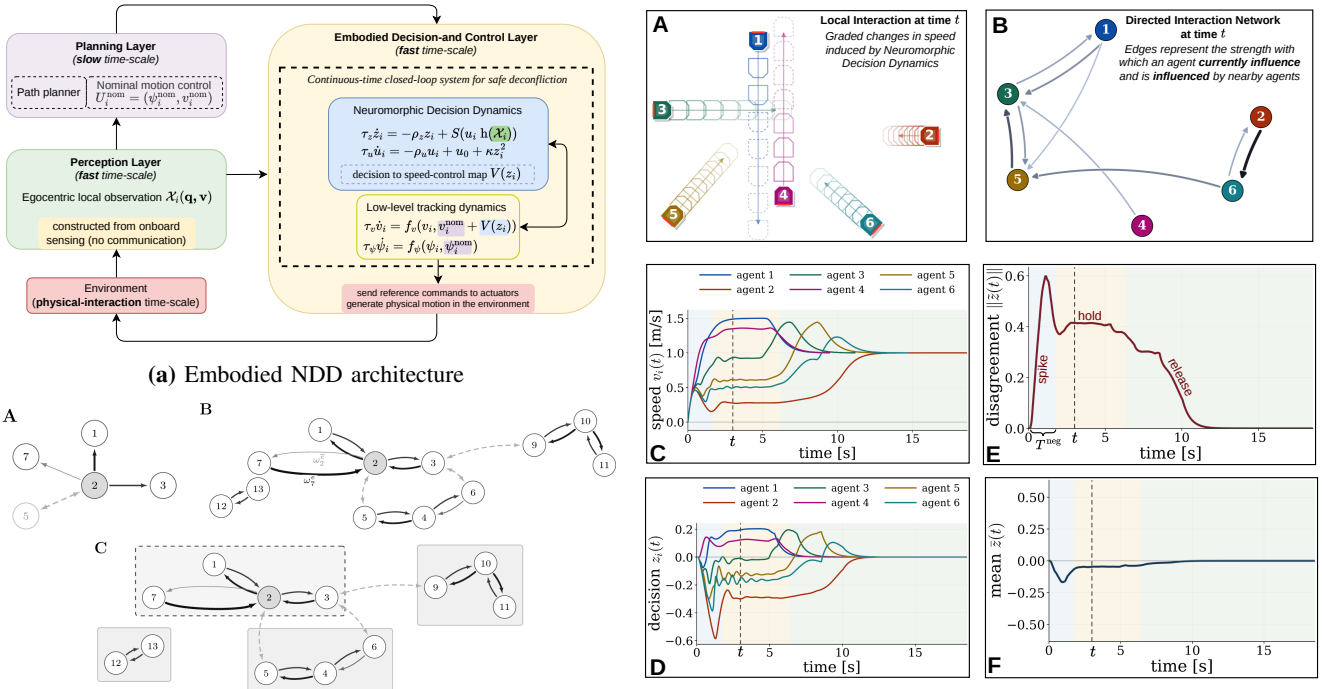
All coefficients are positive: τ_z and τ_u set response times, ρ_z and ρ_u set dissipation, u_0 is the basal drive, and κ_u sets activity-dependent attention. The bounded odd sigmoid S produces a graded decision, and the bounded sign-preserving map V raises speed for $z_i > 0$ and lowers it for $z_i < 0$. Thus, z_i continuously modulates v_i^{cmd}, and negotiation unfolds through continuous-time motion. A low-level controller tracks this command subject to actuation limits. NDD preserves the prescribed path; the slower planner handles steering and static obstacles (Fig. 1(a)). Its nonlinear structure builds on nonlinear opinion dynamics [16], [17].

For each sensed pair $e = (i, j)$, the current velocity directions define projected path rays and a *contested region* within the required clearance d^{safe}. Projected entry and exit times define occupancy intervals: overlap indicates conflict, while separation defines a safe temporal order. Agent i aggregates simultaneous conflicts as $h_i(\mathcal{X}_i) = \sum_{e \in \mathcal{E}_i} P(\mathcal{X}_i^e) W(\mathcal{X}_i^e; \mathcal{X}_i) G(\mathcal{X}_i^e)$. Here $P \geq 0$ measures conflict severity, $W \in [0, 1]$ allocates influence among conflicts, and $G \in [-1, 1]$ encodes the favored order from projected entry-time difference. The function G is positive when the kinematics favor agent i, negative when they favor agent j, and zero at a tie. Separating severity, influence, and direction makes h_i interpretable and tunable. Reciprocal observations generate opposite pairwise preferences without communication.

The agents and physical environment together form a collective recurrent network: $\mathcal{X}_i \mapsto h_i \mapsto (z_i, u_i) \mapsto v_i^{\mathrm{cmd}} \mapsto \mathcal{X}_j$. Intermediate decisions continuously alter motion and neighboring evidence, allowing short-lived cues to propagate through conflict chains without messages or neighbor models. The conflict graph reconfigures as edges appear or disappear and groups merge or split (Fig. 1(b,c)). The guarantees cover these changes. Local parameter conditions tune sensing, attention, speed authority, and response times so the recurrent loop acts before the conflict geometry changes, independently of population size.

B. Bifurcation-driven traversal-order formation

Within a locally engaged group, a group-wide tie makes $\boldsymbol{z} = \mathbf{0}$ a collective-indecision equilibrium with no leading agent. Remaining at this equilibrium produces mutual yielding or freezing—deadlock—while repeatedly changing motion without establishing persistent conflict-edge orientations produces livelock. Let $\boldsymbol{F}$ denote the right-hand side of the group decision dynamics and $\boldsymbol{A} = D_{\boldsymbol{z}} \boldsymbol{h}(\mathcal{X}_{\boldsymbol{z}})$ the sensitivity of group evidence to decision-induced speed changes. Linearization at indecision gives $D_{\boldsymbol{z}} \boldsymbol{F}(\mathbf{0}) = -\rho_z \boldsymbol{I} + S'(0) \frac{u_0}{\rho_u} \boldsymbol{A}$. Collective indecision loses stability when the dominant eigenvalue of $\boldsymbol{A}$ crosses $\lambda^* = \frac{\rho_z \rho_u}{S'(0) u_0}$. We tune conflict pressure and the basal attention level u_0/ρ_u so weak

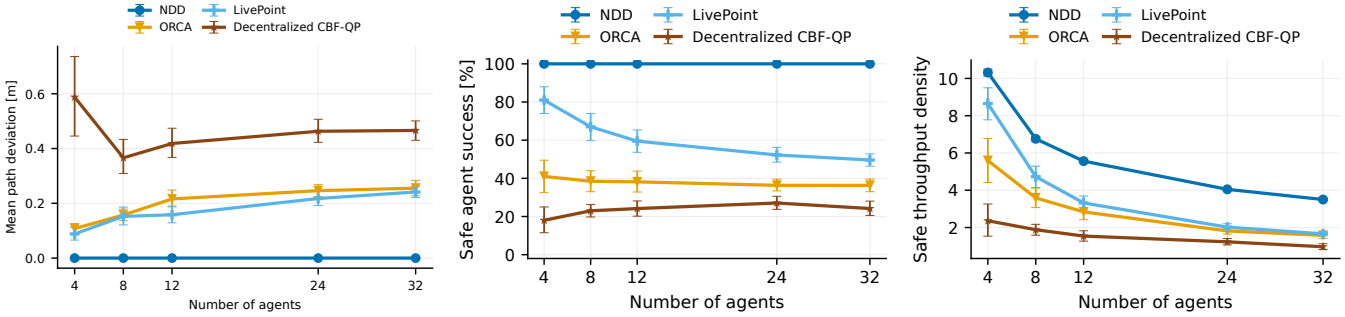


(a) Embodied NDD architecture

(b) Local sensing and conflict-graph hierarchy

(c) Closed-loop safe deconfliction

Fig. 1. **Closed-loop, multi-time-scale NDD architecture and the emergence of safe deconfliction. (a)** A slow planning layer supplies each agent with a prescribed path and nominal speed. At a faster time scale, onboard sensing constructs an egocentric observation from the agent's own motion and the relative motion of nearby agents. Local conflict evidence drives NDD, whose decision state continuously raises or lowers the commanded speed along the prescribed path. The executed motion changes the observed interaction geometry, closing the decision–action–perception loop without agent-to-agent communication. **(b)** Local observations determine which sensed neighbors are in physical conflict and how strongly they influence the agent's decision. (A) shows one agent's sensed interactions, (B) the resulting weighted directed conflict graph, and (C) its weakly connected components. An edge $i \to j$ means that agent j's motion influences agent i's conflict evidence; its weight indicates interaction strength. Each component is a locally engaged group that negotiates independently and concurrently. **(c)** A six-agent example shows the resulting collective response. Internal panels (A) and (B) show agent motion and the active conflict graph; (C) and (D) show realized speeds and decision states, where $z_i > 0$ increases speed and $z_i < 0$ decreases it relative to the nominal speed of $1\,\mathrm{m/s}$; (E) and (F) show decision differentiation and the evolution of the common decision component across the group. NDD amplifies transient motion differences during negotiation, maintains differentiated speeds while a safe traversal order is needed, and returns the agents to nominal motion as conflicts clear. The same closed loop therefore provides collision avoidance (safety) and continued progress (liveness).



(a) Path deviation **(b)** Safe-agent success **(c)** Safe throughput

Fig. 2. **Randomized simulation benchmarks.** (a) Mean cross-track distance from the prescribed path (lower is better). (b) Fraction of agents reaching their goals without a safety violation (higher is better). (c) Violation-free goal-reaching agents per episode duration, normalized by lane count (higher is better). Error bars show variation across the 50 shared scenarios at each group size.

encounters remain below this threshold and safety-critical encounters cross it. The unstable mode contains positive and negative components, so small endogenous asymmetries select a differentiated branch on which some agents accelerate and others yield. The resulting speed differences separate occupancy intervals and orient conflict edges into a persistent traversal order (Fig. 1(c)). As conflicts clear, pressure falls, the graph reconfigures, and decisions return to zero. The tunable time-scale condition $T_{\mathrm{fast}} \lesssim T^{\mathrm{neg}} \ll T_{\mathrm{trav}}$ ensures that decision, attention, and speed tracking integrate short-lived cues within the negotiation horizon, before the contested geometry changes substantially.

C. Safety and liveness guarantees

Under the stated geometric, sensing, tracking, and actuation bounds, our analysis establishes:

1) *Continuous safety:* every conflicting pair maintains $\|\boldsymbol{q}_i - \boldsymbol{q}_j\| \geq d^{\mathrm{safe}}$ through negotiation and graph reconfiguration.
2) *Finite-time order formation:* each fixed locally engaged group forms a strict acyclic order with a robust clearance margin within T^{neg}, and each pairwise order persists until its conflict clears.
3) *Deadlock- and livelock-free progress:* each unresolved group obtains a leading mover and clears a conflict in finite time; repeated clearances resolve the episode while independent groups move concurrently.
4) *Adaptation and release:* conflict removal preserves remaining orders, conflict additions and group mergers renew negotiation, and $(z_i, u_i, v_i) \to (0, u_0/\rho_u, v_i^{\mathrm{nom}})$ after the final conflict clears.

III. Evaluation

We compared NDD with Optimal Reciprocal Collision Avoidance (ORCA) [6], LivePoint [9], and decentralized CBF-QP [7] using homogeneous Dubins-car agents. For each $N \in \{4, 8, 12, 24, 32\}$, all methods received the same 50 randomized start–goal scenarios, giving 250 rollouts per method. Success required reaching the goal without a safety violation. NDD achieved 100% safe-agent success at every group size, preserved the prescribed paths, and produced the highest lane-normalized safe throughput (Fig. 2). LivePoint was the strongest baseline; its safe-agent success rate decreased from approximately 80% at four agents to 50% at 32 agents. ORCA and decentralized CBF-QP remained below 45% across the tested sizes. NDD's longitudinal intervention kept its mean path deviation at zero.

We further evaluated a dense ten-agent near-antipodal exchange. From the same initial states and goals, NDD alone delivered all ten agents without leaving their lanes or violating safety. Its minimum pairwise distance was 1.026 m against a required clearance of 0.7 m; ORCA, LivePoint, and decentralized CBF-QP each incurred at least one safety violation. The outcome illustrates the core mechanism: a group-wide temporal order emerges through continuous speed differentiation even when many pairwise conflicts overlap.

IV. Discussion

NDD establishes a constructive principle for shaping emergent behavior in embodied multi-agent systems. Agents convert local observations into continuous-time evidence, act on it, and use the resulting motion cues to update their decisions. Safety-critical conflict triggers a bifurcation that amplifies small asymmetries into differentiated decisions. This decision–action–perception loop creates temporal separation for collision avoidance and a strict traversal order for progress, then restores nominal motion as conflicts clear. Explicit sensing, actuation, and feedback conditions guarantee finite-time, robust deconfliction under bounded uncertainty and time-varying interaction topologies.

Embedding NDD as a lightweight, fast interaction layer lets slower planners handle static obstacles and direct head-on encounters. Alternative evidence and control maps could extend the framework to road vehicles, UAVs, manipulators, and mixed human–agent teams. Evidence maps could incorporate conventions, intent, or communication, while learning adapts parameters subject to certification conditions. NDD thus provides a general foundation for shaping and certifying collective behavior wherever local physical actions recursively shape collective decisions.

References

[1] J. Grover, C. Liu, and K. Sycara, "Why does symmetry cause deadlocks?" *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 9746–9753, 2020.

[2] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 797–803.

[3] J. van den Berg, M. C. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1928–1935.

[4] Y. Zhou, H. Hu, Y. Liu, S.-W. Lin, and Z. Ding, "A distributed method to avoid higher-order deadlocks in multi-robot systems," *Automatica*, vol. 112, p. 108706, 2020.

[5] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Proceedings of Robotics: Science and Systems*, Ann Arbor, Michigan, Jun. 2016.

[6] J. van den Berg, S. J. Guy, M. C. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research*, ser. Springer Tracts in Advanced Robotics. Springer, 2011, vol. 70, pp. 3–19.

[7] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Transactions on Robotics*, vol. 33, no. 3, pp. 661–674, 2017.

[8] R. Chandra, V. Zinage, E. Bakolas, P. Stone, and J. Biswas, "Deadlock-free, safe, and decentralized multi-robot navigation in social mini-games via discrete-time control barrier functions," *Autonomous Robots*, vol. 49, no. 2, p. 12, 2025.

[9] J. Chen and R. Chandra, "Livepoint: Fully decentralized, safe, deadlock-free multi-robot control in cluttered environments with high-dimensional inputs," *arXiv preprint arXiv:2503.13098*, 2025.

[10] J. Gregoire, S. Bonnabel, and A. de La Fortelle, "Priority-based coordination of robots," *Automatica*, vol. 50, no. 11, pp. 2902–2908, 2014.

[11] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, "Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 1475–1481.

[12] J. E. Hopcroft, J. T. Schwartz, and M. Sharir, "On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the "warehouseman's problem"," *The International Journal of Robotics Research*, vol. 3, no. 4, pp. 76–88, 1984.

[13] J. Yu and S. M. LaValle, "Structure and intractability of optimal multi-robot path planning on graphs," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, vol. 27, 2013, pp. 1443–1449.

[14] M. Jankovic, "Instinctive negotiation by autonomous agents in dense, unstructured traffic: A controls perspective," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2024.

[15] M. Jankovic, M. Santillo, and Y. Wang, "Multiagent systems with CBF-based controllers: collision avoidance and liveness from instability," *IEEE Transactions on Control Systems Technology*, vol. 32, no. 2, pp. 705–712, 2024.

[16] A. Bizyaeva, A. Franci, and N. E. Leonard, "Nonlinear opinion dynamics with tunable sensitivity," *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1415–1430, 2023.

[17] N. E. Leonard, A. Bizyaeva, and A. Franci, "Fast and flexible multiagent decision-making," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, pp. 19–45, 2024.