

The City College of New York

CSC 59867-E: Senior Project II

AI Agents for Decision Making in the Real World

Drone Swarms, Autonomous Vehicles, and Transportation Advanced Topics: Augmented Reality (AR/VR/XR) and Hardware Agents

By Saptarashmi Bandyopadhyay
Email: sbandyopadhyay@ccny.cuny.edu
Assistant Professor of Computer Science
City College of New York and Graduate Center at the City University of New York

October 13, 2026 CSC 59867 - E

Saptarashmi Bandyopadhyay

Lecture Agenda

Distributed Swarm Kinetics and Kinematics

01

- Reynolds/Olfati-Saber flocking control
- Dynamic algebraic graph connectivity (Fiedler value)
- State augmentation under RF transmission latency

Autonomous Transportation and AVs (self-driving cars, drones, ships etc)

02

- Dec-POMDP formulation of unsignalized crossings
- High-speed platoon string stability limits
- Control Barrier Functions (CBFs) for zero-collision safety

AR/VR/XR and Hardware Agents

03

- OpenXR spatial tele-autonomy & retinal HUD vectors
- Sim-to-Real domain randomization (mass, friction, wind)
- Edge deployment under 15W SWaP budgets (Jetson Orin)

Quantum Control vis-a-vis Deterministic AI

04

- QAOA/QUBO combinatorial fleet task allocation
- Parameterized Quantum Circuits (PQCs) for policy search
- Escaping coordination traps via quantum entropy exploration

Saptarshmi Bandyopadhyay

Physical Differences between Discrete Logistics vs. Embodied Fleets

Operational Dimension	Discrete Warehouse Logistics	Embodied Aerial & Highway Fleets
State Space	2D discrete grid, topological rack maps	Continuous $R^3 \times SE(3)$ manifold with non-linear drift
Dynamics & Actuation	Holonomic kinematics; negligible drift	Underactuated, non-holonomic, aerodynamic wake coupling
Communication	Centralized Wi-Fi; deterministic routing	Ad-hoc mesh (MANET); severe packet drop & jitter
Failure Penalty	Stall / operational throughput delay	Physical crash, loss of control, catastrophic damage

Distributed Swarm Kinematics & Flocking Control

Double-Integrator Kinematics

- Kinematic equations: $\dot{p}_i = v_i$, $\dot{v}_i = u_i$ for agent $i \in \{1, ..., N\}$
- u_i synthesizes repulsion, cohesion, and goal navigation
- Gradient potential term: $\nabla\psi(||p_i - p_□||)$ prevents collisions at short range and enforces cohesion at distance
- Actuator velocity saturation limits: $||u_i|| \leq u_max$

Consensus Control Law

- Velocity consensus: $-\sum a_{i□}(t)(v_i - v_□)$ dampens velocity divergence
- Mitigates aerodynamic shear stress and flock turbulence
- Navigation vector f_nav drives centroid toward target waypoints
- Guarantees asymptotic convergence to shared flock velocity

Dynamic Graph Laplacians & State Augmentation for Network Coherence

Bandyopadhyay

Graph Topology & Fiedler Bound

- Dynamic proximity graph $G(t) = (V, E(t))$ with Laplacian $L(t) = D(t) - A(t)$
- Flocking cohesion strictly requires positive algebraic connectivity: $\lambda_2(L(t)) > 0$
- Dynamic edge disconnects occur under RF occlusions and range limits
- Disconnected components trigger split-subswarm divergence

Mitigating Latency via Augmentation

- Observation delay $\tau_i(t)$ breaks standard Markov assumptions
- State augmentation lifts local vector: $\tilde{x}_i(t) = [p_i, v_i, \{p(t-\tau), v(t-\tau)\}]$
- State-space dimensionality scales with communication latency
- Compensates for packet staleness via Extended Kalman Filters (EKF)

Multi-Agent Rollout Planning for Fleet Trajectory Optimization

Bandyopadhyay

Joint Action Explosion

- Joint action space across N continuous agents scales as $O(|U|^N)$
- Exact dynamic programming lookahead is computationally intractable
- Standard MPC re-planning fails to meet real-time loop constraints
- Decentralization is required for scalable embedded execution

Agent-by-Agent Factorization

- Sequential rollout factorizes joint search into N sequential sub-problems
- Each agent minimizes cost given previous choices and baseline policies
- Reduces complexity from $O(|U|^N)$ down to linear $O(|U| \cdot N)$
- Guarantees cost improvement: $J_{rollout}(x) \leq J_{base}(x)$

Connected Self-Driving Autonomous Vehicles (AV)

Saptarashmi Bandyopadhyay

Cooperative AV Systems: Intersections vs. High-Speed Platooning

Unsignalized Intersections

- V2V/V2I peer coordination replaces static traffic signals
- Spatio-temporal reservation grids prevent crossing collisions
- Asymmetric traffic arrivals cause priority deadlock traps
- Blind spots and building occlusions restrict local visibility

Highway Platooning

- Close-following platoons (<1m gap) cut aerodynamic drag by 25%
- String stability condition: $||y_i(s) / y_{i-1}(s)|| \leq 1$
- Braking disturbances must attenuate upstream along the convoy
- V2V packet delay jitter triggers shockwave crashes without feedforward

Bandyopadhyay

Sensory Perception Architecture: 3D Meshes, BEV & V2X Streams

Multi-Modal Sensor Suite

- Sparse LiDAR point clouds processed via PointNet / VoxelNet
- Surround-view camera streams projected to Bird's-Eye-View (BEV)
- V2X edge messaging ingests neighboring velocities and brake intents
- Closed-loop simulators: CARLA, Isaac Sim, AirSim linked via ROS2

Local Observation Tuple

- Observation $o_i = [p_i, v_i, \theta_i, \omega_i, BEV_map, \{(p_j - p_i, v_j - v_i)\}]$
- Egocentric coordinate normalization preserves spatial translation invariance
- Temporal frame stacking tracks vehicle accelerations under occlusion
- Graph Neural Networks (GNNs) aggregate dynamic V2V neighbor edges

Failure Modes of Classical Transportation Heuristics

Time-To-Collision (TTC)

- Conservative threshold alerts (TTC < τ_crit)
- Freezing Robot Problem: vehicle halts indefinitely in dense merge traffic
- Static margins cannot anticipate collaborative yielding

Static Right-of-Way

- Rigid priority hierarchies (FIFO, Main-Road Priority)
- Causes chronic starvation and gridlock on secondary approaches
- Lacks dynamic responsiveness to asymmetric traffic bursts

Fixed Distance Spacing

- Classical PID gap regulation without feedforward V2V signals
- Exhibits string instability across long vehicle platoons
- Braking oscillations amplify upstream into phantom traffic jams

Fleet Orchestration Formulated as a Dec-POMDP / POSG

Formal Tuple G

- $G = (I, S, \{A^i\}, \{O^i\}, P, \Omega, \{R^i\})$ for $i \in \{1, ..., N\}$
- State S: Joint poses, velocities, and road occupancy maps
- Partial observation O^i: Ego sensor field bounded by range and obstacles
- Transition P: Real-world vehicle kinematics subject to tire slip

Multi-Objective Reward Structure

- Progress reward: $w_v \cdot v_{progress}$ encourages route completion
- Comfort penalties: $-w_{jerk} \cdot ||\ddot{v}||^2$ minimizes passenger motion sickness
- Energy regularization: $-w_e \cdot ||u||^2$ penalizes aggressive actuation
- Safety penalty: Large terminal penalty on collision or critical boundary violation

Questions?

Saptarashmi Bandyopadhyay