Hierarchical Neuro-Semantic Control for UAV Swarm Formation: Bridging LLM Planning and Hard Safety Constraints

1st Xianchang Wang*
Shenyang Aerospace University
School of Computer Science
Shenyang, China
wxc@email.sau.edu.cn

2nd Jie Sun
Shenyang Aerospace University
School of Computer Science
Shenyang, China
sunjie8742@stu.sau.edu.cn

***Abstract*—The integration of high-level semantic planning with low-level physical constraints remains a critical challenge for UAV swarm formation. While Large Language Models (LLMs) demonstrate proficiency in task reasoning, their inherent non-determinism and inference latency impede their utilisation in safety-critical systems. To address this issue, we propose Hierarchical Neuro-Semantic Control (HNSC), an architecture that achieves non-authoritative semantic planning with conditional hard safety via HO-CBF-QP filtering when the QP remains feasible. At the semantic layer, a contract-based logical framework constrains ambiguous intentions into verifiable command sets. At the Translation layer, a robust High-Order Control Barrier Function (HO-CBF), coordinated with a Reference Governor, performs real-time safety filtering under measurement uncertainties. Moreover, the system incorporates a deadlock monitor that triggers strategy reconstruction when the swarm encounters local minima. Experiments on the `gym-pybullet-drones` platform show that, in a challenging narrow-gap traversal task, HNSC achieves a 95.4% success rate and reduces the deadlock rate from 41.0% to 1.8% compared with open-loop LLM baselines. Collision events are significantly reduced by HO-CBF-QP filtering during feasible QP execution, while sampled distance statistics remain close to the execution boundary under discrete-time reporting. Overall, the proposed architecture improves swarm adaptability and interpretability while providing conditional hard-safety guarantees through HO-CBF-QP filtering.**

***Index Terms*—Large Language Models; UAV Swarm; Formation Control; High-Order Control Barrier Functions; Quadratic Programming; Robust Safety**

I. Introduction

Multi-UAV swarms show strong potential in missions such as search and rescue [1]–[3], but their deployment remains limited by a key mismatch: high-level tasks are semantic and ambiguous, whereas low-level execution must satisfy strict geometric safety constraints.

Classical formation control [4] and collision avoidance methods [5] provide rigorous coordination mechanisms, yet lack the semantic grounding needed for natural-language task specification. In contrast, LLMs exhibit strong capabilities in zero-shot planning [6]–[8] and embodied reasoning [9], [10], but their non-determinism, hallucinations, and inference latency [11] make them unsuitable as authoritative components in safety-critical control. Meanwhile, Control Barrier Functions (CBFs) [12]–[14], including High-Order CBFs [15], measurement-robust variants [16], [17], and Reference Governors [18], [19], can enforce safety constraints, but simply cascading them with an LLM is insufficient: open-loop semantic plans may still conflict with environmental geometry and trap the swarm in safe but non-progressing local minima.

To address this gap, we propose HNSC, following the principle of decoupling non-authoritative semantic reasoning from execution, where a robust HO-CBF-QP filter enforces conditional hard safety as long as the QP remains feasible. HNSC constrains LLM outputs through a contract-based interface [20], enforces hard safety with a robust HO-CBF-QP filter, and closes the loop with deadlock monitoring that triggers semantic strategy reconstruction when progress stalls. We evaluate the framework on `gym-pybullet-drones` [21], [22] and show that this closed-loop design substantially improves reachability in strongly constrained environments while preserving safety.

The main contributions are:

- We propose a hierarchical asynchronous semantic-execution closed-loop architecture for UAV swarms, coupling LLM-based planning with hard safety constraints in a composable manner.
- We design a contract-based semantic planner and structured interface for verifiable LLM outputs, improving interpretability and execution reliability.
- We develop a measurement-robust HO-CBF with a Reference Governor to transform high-relative-degree safety constraints into QP-solvable inequalities for real-time "hard-safety/soft-performance" filtering.
- We introduce a deadlock monitor with signal-triggered strategy reconstruction, enabling recovery from safe-but-stuck local minima in strongly constrained environments.

II. Related Work

Multi-UAV formation and consensus control provide scalable coordination primitives, but dense obstacles require explicit handling of collision and input constraints. CBF-QP

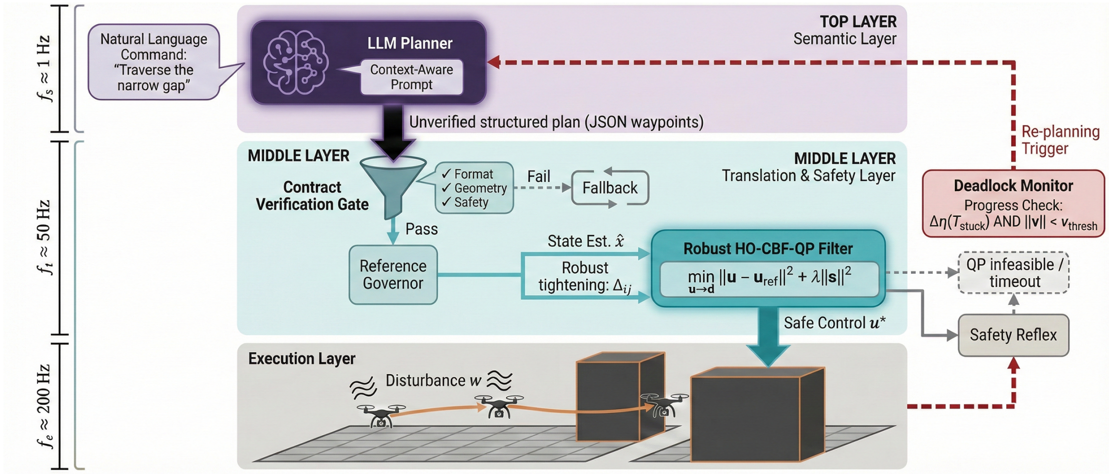

Fig. 1: Overview of HNSC: hierarchical asynchronous neuro-semantic closed loop.

safety filters and safety certificates enforce forward-invariant safe sets for multi-robot collision avoidance. High relative-degree constraints in quadrotors motivate HO-/exponential CBFs, and measurement-robust formulations address sensing errors and disturbances. LLM-based robot planning improves semantic reasoning, but prior approaches typically lack verifiable hard safety guarantees for coupled multi-UAV constraints. We leverage contract-based design to bound and verify semantic outputs before execution.

III. Methodology

A. Preliminaries and Problem Formulation

We consider a hierarchical closed-loop architecture that reconciles non-deterministic LLM planning with hard physical constraints in UAV swarms.

a) Cascaded Dynamics: To balance modeling fidelity and real-time optimization, the outer loop models the i-th UAV as a disturbed double integrator:

$$\dot{\mathbf{p}}_i = \mathbf{v}_i, \quad \dot{\mathbf{v}}_i = \mathbf{u}_i + \mathbf{w}_i \tag{1}$$

where $\mathbf{p}_i, \mathbf{v}_i \in \mathbb{R}^3$ denote position and velocity, $\mathbf{u}_i \in \mathbb{R}^3$ is the nominal outer-loop acceleration, and $\mathbf{w}_i \in \mathcal{W}$ is a bounded disturbance capturing unmodeled dynamics, aerodynamic effects, parameter uncertainty, and inner-loop tracking error.

To ensure dynamic feasibility, the control input is restricted to the achievable acceleration set:

$$\mathbf{u}_i \in \mathcal{U}_a := \left\{ \mathbf{u} \in \mathbb{R}^3 : \|\mathbf{u}\|_\infty \le u_{\max} \right\} \tag{2}$$

where $u_{\max}$ is determined by the platform thrust and tilt limits.

b) Safety Set and Robust Sensing: We define the global safe set using the minimum separation threshold $d_{\min}$ for both inter-UAV and UAV–obstacle interactions. For any UAV pair (i, j), let

$$h_{ij}(\mathbf{x}) := \|\mathbf{p}_i - \mathbf{p}_j\|^2 - d_{\min}^2 \tag{3}$$

and for any static obstacle O_k, define

$$h_{ik}^{\text{obs}}(\mathbf{p}_i) := \text{dist}(\mathbf{p}_i, O_k)^2 - d_{\min}^2. \tag{4}$$

The nominal safe set is then

$$\begin{aligned} \mathcal{S} := \big\{ \mathbf{x} \mid\, & h_{ij}(\mathbf{x}) \ge 0,\ \forall (i,j), \\ & \text{and } h_{ik}^{\text{obs}}(\mathbf{p}_i) \ge 0,\ \forall (i,k) \big\}. \end{aligned} \tag{5}$$

Thus, safety requires both inter-agent and UAV–obstacle distances to remain no smaller than $d_{\min}$, which later yields non-relaxable HO-CBF constraints in the QP.

To account for perception and communication uncertainty, we assume bounded disturbance and bounded state-estimation error: *Assumption 1:* $\|\mathbf{w}_i\| \le \bar{w}$ for all i and all time. *Assumption 2:*

$$\mathbf{p}_i = \hat{\mathbf{p}}_i + \mathbf{e}_i^p, \quad \mathbf{v}_i = \hat{\mathbf{v}}_i + \mathbf{e}_i^v \tag{6}$$

with $\|\mathbf{e}_i^p\| \le \epsilon_p$ and $\|\mathbf{e}_i^v\| \le \epsilon_v$. For pairwise safety analysis, the bounds $\bar{w}, \epsilon_p, \epsilon_v$ used below are understood as conservative worst-case bounds after relative aggregation, i.e., on $\|\mathbf{w}_{ij}\|$, $\|\mathbf{e}_i^p - \mathbf{e}_j^p\|$, and $\|\mathbf{e}_i^v - \mathbf{e}_j^v\|$.

Based on these bounds, we define a conservative robust safe domain by shrinking $\mathcal{S}$ via the Minkowski difference:

$$\mathcal{S}^{\text{rob}} := \mathcal{S} \ominus \mathbb{B}(0, \delta), \tag{7}$$

where δ upper-bounds the worst-case deviation caused by estimation error and delay-induced drift. Equivalently, the implementation enforces

$$\|\hat{\mathbf{p}}_i - \hat{\mathbf{p}}_j\| \ge d_{\min} + \delta, \qquad \text{dist}(\hat{\mathbf{p}}_i, O_k) \ge d_{\min} + \delta. \tag{8}$$

Accordingly, we define the execution-layer hard safety threshold

$$d_{\text{safe}} := d_{\min} + \delta. \tag{9}$$

In practice, we choose

$$\delta = \epsilon_p + v_{\max}\tau, \tag{10}$$

where τ is the worst-case sensing/communication delay and $v_{\max}$ is a known velocity bound. In experiments, $d_{\min} = 0.30$ m, and d_{safe} is computed from (10).

B. Hierarchical Neuro-Semantic Control Framework

As shown in Fig. 1, HNSC adopts an asynchronous semantic-execution dual-loop architecture that decouples high-level intent reasoning from low-level safety projection, thereby preserving safe degradation when the semantic layer fails.

a) Formalized Semantic Planner via LLM: The semantic layer outputs *structured waypoint plans* instead of free-form text. At each semantic update, the LLM maps the observation $\mathbf{o}_t$—including active UAV IDs, role assignments, mission progress, and anomaly flags ξ_t (e.g., STUCK/BLOCKED)—to a JSON waypoint plan Π_t. The plan contains `reasoning`, `waypoints`, and optional `new_roles`, where `waypoints` assigns each active UAV a current 3D waypoint $\mathbf{g}_i$ (or the first element of a short verified sequence) over the next planning horizon. This makes the semantic layer *non-authoritative*: its output is executed only after deterministic verification and feasibility checking.

Semantic Output Contract. A semantic output is accepted only if it: (i) is valid JSON with required fields; (ii) assigns exactly one current waypoint to each active UAV with consistent IDs; (iii) keeps all waypoints within workspace bounds $[-4.5, 4.5]^2 \times [0.3, 3.0]$; (iv) places all waypoints in free space, i.e., outside static obstacles and their planning-time inflated geometry; and (v) satisfies the semantic-level separation constraint $\|\mathbf{g}_i - \mathbf{g}_j\| \geq d_{\min}$ for all $i \neq j$. Otherwise, it is rejected by the verification gate. The semantic contract uses the nominal threshold $d_{\min}$ for coarse plan screening, while the Translation layer enforces the tightened hard threshold $d_{\text{safe}} = d_{\min} + \delta$ during execution.

b) Asynchronous Bridge and Verification Gate: To isolate LLM inference latency from the control loop, the framework introduces a feasibility probe, asynchronous degradation, and strategy hysteresis. Given a candidate plan, a nominal reference $\mathbf{u}^{\text{ref}}$ toward $\mathbf{g}_i$ is constructed and checked by a one-step QP feasibility test under input bounds and current hard HO-CBF constraints; infeasible plans are rejected and replaced by a fallback strategy. If the semantic layer times out or violates the contract, the Translation layer retains the last verified strategy and falls back to a deterministic backup planner. To avoid oscillatory switching, a new semantic plan is accepted only when the elapsed time since the last update exceeds the hysteresis threshold T_{hyst}.

C. Unified Safety via Robust High-Order CBF-QP

To enforce hard safety under relative-degree-2 collision constraints with estimation errors and disturbances, we adopt a robust HO-CBF formulation at the Translation layer.

a) Measurement-Robust High-Order CBF Formulation: For any UAV pair (i, j), define

$$h_{ij}(\mathbf{p}) = \|\mathbf{p}_i - \mathbf{p}_j\|^2 - d_{\text{safe}}^2 \tag{11}$$

with safety requirement $h_{ij} \geq 0$, where $d_{\text{safe}} = d_{\min} + \delta$ is defined in (10). For a static obstacle O_k, the same construction is applied to the distance function $h_{ik}^{\text{obs}}(\mathbf{p}_i) = \text{dist}(\mathbf{p}_i, O_k)^2 - d_{\text{safe}}^2$; below we write the UAV–UAV case explicitly, while the obstacle case is specialized with zero obstacle velocity and acceleration.

Under (1), let $\mathbf{p}_{ij} = \mathbf{p}_i - \mathbf{p}_j$, $\mathbf{v}_{ij} = \mathbf{v}_i - \mathbf{v}_j$, $\mathbf{u}_{ij} = \mathbf{u}_i - \mathbf{u}_j$, and $\mathbf{w}_{ij} = \mathbf{w}_i - \mathbf{w}_j$. Then

$$\begin{aligned} \dot{h}_{ij} &= 2\mathbf{p}_{ij}^\top \mathbf{v}_{ij} \\ \ddot{h}_{ij} &= 2\mathbf{v}_{ij}^\top \mathbf{v}_{ij} + 2\mathbf{p}_{ij}^\top (\mathbf{u}_{ij} + \mathbf{w}_{ij}) \end{aligned} \tag{12}$$

and the second-order HO-CBF condition is

$$\ddot{h}_{ij} + k_1 \dot{h}_{ij} + k_0 h_{ij} \geq -\Delta_{ij}, \quad k_0, k_1 > 0 \tag{13}$$

where k_0, k_1 are chosen such that $\lambda^2 + k_1\lambda + k_0$ is Hurwitz. Substituting (12) into (13) yields the affine control constraint

$$2\mathbf{p}_{ij}^\top \mathbf{u}_{ij} \geq \hat{b}_{ij} - \Delta_{ij} \tag{14}$$

where Δ_{ij} upper-bounds the effect of bounded disturbances and estimation errors. In implementation, we use

$$\begin{aligned} \Delta_{ij} = {} & 2\|\hat{\mathbf{p}}_{ij}\|\bar{w} + 2\|\hat{\mathbf{v}}_{ij}\|\epsilon_v \\ & + k_1(2\|\hat{\mathbf{p}}_{ij}\|\epsilon_v + 2\|\hat{\mathbf{v}}_{ij}\|\epsilon_p) + k_0(2\|\hat{\mathbf{p}}_{ij}\|\epsilon_p) \end{aligned} \tag{15}$$

which is evaluated online (or conservatively pre-bounded) using $\bar{w}, \epsilon_p, \epsilon_v$. The tightened term is

$$\hat{b}_{ij} := -2\hat{\mathbf{v}}_{ij}^\top \hat{\mathbf{v}}_{ij} - k_1\left(2\hat{\mathbf{p}}_{ij}^\top \hat{\mathbf{v}}_{ij}\right) - k_0\left(\|\hat{\mathbf{p}}_{ij}\|^2 - d_{\text{safe}}^2\right) \tag{16}$$

so that (14) can be directly incorporated into the QP as a tightened safety constraint.

b) Soft-Performance QP: At 50 Hz, the Translation layer projects the semantic reference $\mathbf{u}^{\text{ref}}$ onto the set satisfying hard HO-CBF safety constraints and input bounds. Slack is introduced only for performance terms, while safety constraints remain non-relaxable. Let $\mathbf{u} = [\mathbf{u}_1^\top, \ldots, \mathbf{u}_N^\top]^\top$. Stacking (14) yields $\mathbf{A}_{\text{cbf}}(\mathbf{x})\mathbf{u} \geq \mathbf{b}_{\text{cbf}}(\mathbf{x})$, and performance inequalities are written as $\mathbf{A}_{\text{perf}}(\mathbf{x})\mathbf{u} \geq \mathbf{b}_{\text{perf}}(\mathbf{x}) - \boldsymbol{s}$, $\boldsymbol{s} \geq 0$. The QP is

$$\begin{aligned} \min_{\mathbf{u}, \boldsymbol{s}} \quad & \left\|\mathbf{u} - \mathbf{u}^{\text{ref}}\right\|_Q^2 + \lambda \left\|\boldsymbol{s}\right\|_2^2 \\ \text{s.t.} \quad & \mathbf{A}_{\text{cbf}}(\mathbf{x})\mathbf{u} \geq \mathbf{b}_{\text{cbf}}(\mathbf{x}) \\ & \mathbf{A}_{\text{perf}}(\mathbf{x})\mathbf{u} \geq \mathbf{b}_{\text{perf}}(\mathbf{x}) - \boldsymbol{s} \\ & \boldsymbol{s} \geq 0 \\ & \|\mathbf{u}_i\|_\infty \leq u_{\max}, \quad \forall i \in \mathcal{V}. \end{aligned} \tag{17}$$

Proposition 1 (Conditional Safety): If initially $h_{ij}\big(\mathbf{x}(0)\big) \geq 0$ for all UAV pairs (i, j) and $h_{ik}^{\text{obs}}\big(\mathbf{p}_i(0)\big) \geq 0$ for all UAV–obstacle pairs (i, k), and the QP remains feasible while enforcing the hard HO-CBF constraints at all times, then the HO-CBF safe set is forward invariant during feasible QP execution. The UAV–obstacle case follows analogously through the specialized static-obstacle distance function.

If the QP times out or becomes infeasible, the controller switches to the safety reflex

$$\mathbf{u}_i = \text{sat}_{u_{\max}}\left(-k_b\mathbf{v}_i + \mathbf{u}_i^{\text{repel}}\right) \tag{18}$$

where $\mathbf{u}_i^{\text{repel}}$ is a small repulsive term based on nearest-neighbor or obstacle direction. This mode prioritizes risk mitigation rather than task optimality.

D. Deadlock Resolution and Decentralized Contingency

Because local QP-based methods can stall in complex geometric constraints, the system may enter local minima or deadlock states. We therefore introduce a progress-based monitor and a decentralized contingency mechanism to preserve survivability when progress stagnates or the semantic layer fails.

For each agent, task progress is measured by the distance to its semantic target, $\eta_i(t) = \|\mathbf{p}_i(t) - \mathbf{g}_i\|$. Deadlock is detected when

$$\mathcal{D}_i := \{\eta_i(t - T_{\text{stuck}}) - \eta_i(t) < \epsilon_{\text{prog}}\} \cap \{\|\mathbf{v}_i\| < v_{\text{thresh}}\} \quad (19)$$

holds for longer than T_{stuck}. In that case, the anomaly flag $\tilde{\xi}_i$ is set to STUCK or BLOCKED and fed back to the semantic layer for strategy reconstruction. To avoid oscillatory updates, a cooling interval T_{cool} is enforced between consecutive strategy changes.

If the semantic layer becomes unavailable due to communication loss or repeated replanning failure, the system enters a "Best-Effort Hazard Mitigation" mode that prioritizes safety over task optimality. Specifically, it applies coordinated braking $\mathbf{u}_i^{\text{brake}} = -k_b \mathbf{v}_i$, adds a local repulsive term $\mathbf{u}_i^{\text{repel}}$ to increase inter-agent spacing, and performs conditional landing or hover when locally safe.

E. Overall Asynchronous Control Loop

The system uses a tri-rate asynchronous architecture to isolate LLM latency from flight-control timing: the semantic layer updates at $f_s \approx 1\,\text{Hz}$, the Translation layer performs monitoring and safety projection at $f_t \approx 50\,\text{Hz}$, and the execution layer tracks thrust/attitude commands at $f_e \approx 200\,\text{Hz}$.

At each Translation-layer step, the controller updates state estimates, checks deadlock flags, accepts a new semantic plan only if it passes contract and feasibility checks, and then applies the Reference Governor and HO-CBF-QP projection before sending commands to the inner loop.

IV. Experiments and Results

We evaluate HNSC on `gym-pybullet-drones` to assess its reliability and adaptability in strongly constrained environments. Although we focus on a representative narrow-gap scenario, the goal is to validate the proposed closed-loop mechanism rather than optimize a scenario-specific planner. The evaluation targets three aspects: hard safety, task reachability under tight geometric constraints, and deadlock resolution in safe-but-stuck situations.

A. Experimental Setup

Simulator and control rates: We use gym-pybullet-drones [21] with the PyBullet physics engine [22]. The system runs in a tri-rate asynchronous loop: semantic planning at $f_s \approx 1\,\text{Hz}$, Translation-layer monitoring and safety projection at $f_t \approx 50\,\text{Hz}$, and low-level execution at $f_e \approx 200\,\text{Hz}$.

Scenario: Narrow Gap. The main scenario requires $N = 3$ UAVs to move from the start region to the goal region through a $0.8\,\text{m}$-wide doorway. This setting creates a narrow feasible domain in which local-optimal control and coupled collision avoidance easily induce stagnation, making it a representative stress test for deadlock detection and strategy reconstruction.

Safety constraints. The nominal minimum separation is set to $d_{\min} = 0.30\,\text{m}$ for both inter-UAV and UAV–obstacle interactions. The Translation-layer HO-CBF-QP enforces the conservative hard threshold $d_{\text{safe}} = d_{\min} + \delta$, with δ defined in (10), and treats HO-CBF inequalities as non-relaxable constraints; slack is introduced only for performance terms. Collision Rate is computed from physics-engine contact events, whereas reported minimum distances are sampled diagnostics from discretely logged simulator states. In all experiments, $\epsilon_p = 0.01\,\text{m}$, $v_{\max} = 1.0\,\text{m/s}$, and $\tau = 0.04\,\text{s}$.

Deadlock and semantic planning. Deadlock is detected by the sliding-window progress criterion in (19) and triggers event-driven replanning. The semantic planner runs asynchronously with a low-temperature LLM setting to reduce output variance.

Baselines:

All methods share the same low-level execution stack; when enabled, the same Translation-layer HO-CBF-QP is used for fair comparison.

- **Baseline 1: Local Reactive (No-LLM + No CBF-QP).** Local reactive APF control, without semantic planning or hard safety filtering.
- **Baseline 2: Open-loop LLM (with CBF-QP).** One-shot LLM strategy with HO-CBF-QP, but without state-triggered closed-loop replanning.
- **Baseline 3: Classical Planner + Tracking (No-LLM) with CBF-QP.** Geometry-aware A* + spline tracking with HO-CBF-QP, but without semantic reasoning.
- **HNSC (Ours).** Full semantic-physical closed loop with contract verification, robust HO-CBF-QP, and deadlock-triggered replanning.

B. Quantitative Performance Comparison

We evaluate task completion and safety over 72 trials (mean $\pm$ std). Success Rate, Collision Rate, and Safe Deadlock Rate quantify task outcome; $d_{\min}^{\text{UAV}}$ and $d_{\min}^{\text{obs}}$ are sampled near-boundary diagnostics only, because hard safety is enforced at $d_{\text{safe}} = d_{\min} + \delta$ in the $50\,\text{Hz}$ Translation loop whereas collisions are counted from physics-engine contact events.

Table I separates two effects: HO-CBF-QP mainly suppresses collisions (79.8% for Baseline 1 versus 3.2%/3.0% for Baselines 2/3), whereas semantic closed-loop replanning mainly reduces safe deadlock and improves reachability (41.0% for Baseline 2 versus 1.8% for HNSC). HNSC therefore improves success by escaping geometry-induced local minima rather than by relaxing safety margins.

Because hard safety is enforced at $50\,\text{Hz}$ from estimated states while reported minima are computed from discretely logged simulator states and contact-distance queries, sampled minima may occasionally approach or slightly fall below $d_{\min}$ without contact; we therefore treat Collision Rate as the primary safety indicator and use sampled minima only as boundary-tight diagnostics.

TABLE I: Performance on Narrow Gap

Method	Success Rate ↑	Collision Rate ↓	Deadlock Rate ↓	$d_{\min}^{\text{UAV}}$ (m) ↑	$d_{\min}^{\text{obs}}$ (m) ↑
Baseline 1	$15.3\% \pm 4.2\%$	$79.8\% \pm 4.7\%$	$4.9\% \pm 2.5\%$	0.08 ± 0.12	0.05 ± 0.09
Baseline 2	$55.6\% \pm 5.8\%$	$3.2\% \pm 2.1\%$	$41.0\% \pm 5.8\%$	0.31 ± 0.04	0.29 ± 0.05
Baseline 3	$78.0\% \pm 4.9\%$	$3.0\% \pm 2.0\%$	$19.0\% \pm 4.6\%$	0.32 ± 0.03	0.31 ± 0.04
Ours	$\mathbf{95.4\% \pm 2.4\%}$	$\mathbf{2.8\% \pm 1.9\%}$	$\mathbf{1.8\% \pm 1.6\%}$	$\mathbf{0.33 \pm 0.03}$	$\mathbf{0.32 \pm 0.04}$

TABLE II: Ablation Study

Variant	Success Rate ↑	Collision Rate ↓	Deadlock Rate ↓	$d_{\min}^{\text{UAV}}$ (m) ↑	$d_{\min}^{\text{obs}}$ (m) ↑	Δ Success ↑
HNSC Full	95.4%	2.8%	1.8%	0.33	0.32	—
w/o Semantic Closed-Loop	55.8%	3.2%	41.0%	0.31	0.29	-39.6%
w/o Deadlock Monitor	71.5%	3.0%	25.5%	0.32	0.30	-23.9%
w/o HO-CBF-QP	74.3%	19.7%	6.0%	0.12	0.09	-21.1%
w/o Contract Verify	85.2%	3.5%	11.3%	0.30	0.28	-10.2%
w/o Safety Reflex	89.8%	4.5%	5.7%	0.29	0.27	-5.6%

TABLE III: Safety Metrics by Mission Phase

Mission Phase	Duration	$d_{\min}^{\text{UAV}}$ (m)	$d_{\min}^{\text{obs}}$ (m)	CBF Active (%)
Takeoff & Formation	0-5s	0.45 ± 0.08	N/A	12.3%
Approach	5-12s	0.38 ± 0.05	0.41 ± 0.06	35.7%
Narrow Gap Transit	12-18s	$\mathbf{0.33 \pm 0.03}$	$\mathbf{0.32 \pm 0.04}$	**78.4%**
Exit & Recovery	18-25s	0.42 ± 0.06	0.48 ± 0.07	18.2%

TABLE IV: Computational Performance Analysis

Component	Mean Time (ms)	Max Time (ms)	Frequency
State Estimation	0.3	0.8	50 Hz
Deadlock Monitor	0.1	0.3	50 Hz
Reference Governor	0.5	1.2	50 Hz
HO-CBF-QP Solving	**2.1**	**8.5**	50 Hz
LLM Inference	950	1800	∼1 Hz

Table II shows a functional separation: the semantic closed loop is most critical for reachability and deadlock recovery (55.8% success and 41.0% deadlock without it), whereas HO-CBF-QP is the primary hard-safety mechanism (19.7% collision and 0.12/0.09 m minima without it). The deadlock monitor, contract verification, and safety reflex further improve robustness by enabling timely replanning and filtering aggressive plans.

Phase-wise statistics in Table III show that the narrow-gap transit phase is the most safety-critical segment, with the smallest sampled distances and the highest CBF activation rate (78.4%). Computational results in Table IV indicate that HO-CBF-QP solving remains compatible with the 50 Hz Translation loop (2.1 ms mean, 8.5 ms max), while LLM inference operates asynchronously at approximately 1 Hz.

C. Deadlock Resolution in Narrow Gap Traversal

To illustrate the interplay between physical constraints and semantic adaptation, we present a representative "deadlock in front of door" case. Fig. 2 visualizes the evolution process in simulation.

Phase I – Saturation (T = 10–14 s): The swarm approaches the narrow doorway in a triangular formation. As the UAVs enter the constrained region, goal attraction and collision avoidance become conflicting, shrinking the feasible set for

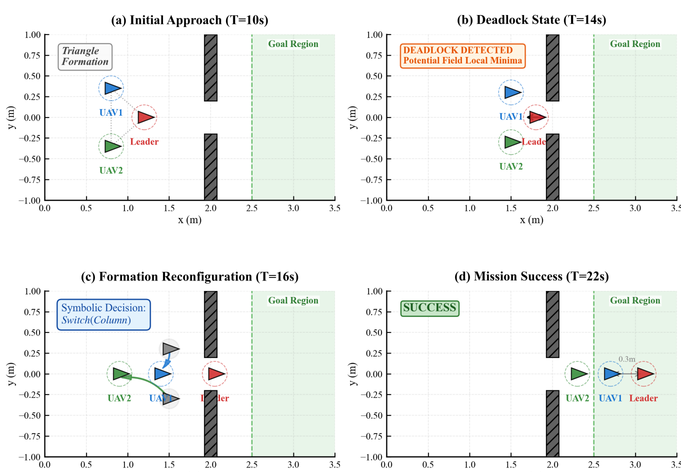

Fig. 2: Snapshot sequence of narrow gap traversal demonstrating deadlock detection and semantic replanning.
(a) Initial triangle formation approaching the gap. (b) Deadlock state: opposing forces cause velocity saturation. (c) LLM-triggered formation reconstruction to column. (d) Successful traversal maintaining safety distances.

motion. Consequently, the CBF-QP increasingly modifies the semantic reference to satisfy the HO-CBF constraints, and the swarm exhibits near-zero velocities with small-amplitude oscillations around the doorway.

Phase II – Triggering (T = 14–15 s): The deadlock monitor detects insufficient progress and low speed, satisfying Equation 19. After persisting for $T_{\text{stuck}} = 3s$, the system raises the STUCK flag and triggers semantic replanning. Notably, conditional hard safety is preserved throughout the detection interval as long as the HO-CBF-QP remains feasible and the hard constraints are enforced at each control step.

Phase III – Adaptation (T = 15–22 s): Upon receiving the deadlock flag and the current formation state, the semantic layer attributes the stagnation to a formation–passage mismatch and outputs a revised strategy (e.g., switching to a single-file/column formation with a sequential passage order). The Reference Governor validates the candidate plan via a one-step QP feasibility test under input bounds and current hard HO-CBF constraints, and the Translation layer executes the transformation while the HO-CBF-QP maintains safety. In this case, inter-UAV distances remain above 0.33 m during the maneuver.

Table III reports phase-wise safety statistics. The narrow-gap transit segment exhibits the highest CBF activation rate (78.4%), indicating the most severe constraint stress, while collision events remain rare and the sampled minimum distances stay close to the execution boundary (with the Translation-layer HO-CBF-QP enforcing $d_{\text{safe}} = d_{\min} + \delta$).

V. Conclusion

This paper proposes HNSC, a hierarchical asynchronous semantic-execution closed-loop architecture for UAV swarm formation tasks under complex geometric constraints. Following the principle of "non-authoritative semantic planning with conditional hard safety via HO-CBF-QP filtering when the QP remains feasible", the framework constrains LLM-generated high-level strategies within a verifiable contract interface and projects semantic outputs to safe executable commands through a robust HO-CBF-QP layer, without relying on the semantic model as an authoritative safety component. In addition, deadlock monitoring and signal-triggered strategy reconstruction enable recovery from safe-but-stuck local minima in strongly constrained environments, improving reachability and overall reliability. We focus on static geometric constraints with $N = 3$; future work will extend to dynamic obstacles, larger swarms, and real-world sensing/communication impairments.

References

[1] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," *IEEE Transactions on automatic control*, vol. 51, no. 3, pp. 401–420, 2006.

[2] 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.

[3] W. Ren and R. W. Beard, *Distributed consensus in multi-vehicle cooperative control: theory and applications*. Springer, 2008.

[4] E. Soria, F. Schiano, and D. Floreano, "Predictive control of aerial swarms in cluttered environments," *Nature Machine Intelligence*, vol. 3, no. 6, pp. 545–554, 2021.

[5] J. Van den Berg, M. 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.

[6] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, "Language models as zero-shot planners: Extracting actionable knowledge for embodied agents," in *International conference on machine learning*. PMLR, 2022, pp. 9118–9147.

[7] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," *arXiv preprint arXiv:2204.01691*, 2022.

[8] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar *et al.*, "Inner monologue: Embodied reasoning through planning with language models," *arXiv preprint arXiv:2207.05608*, 2022.

[9] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, A. Wahid, J. Tompson, Q. Vuong, T. Yu, W. Huang *et al.*, "Palm-e: An embodied multimodal language model," 2023.

[10] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid *et al.*, "Rt-2: Vision-language-action models transfer web knowledge to robotic control," in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.

[11] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. R. Narasimhan, and Y. Cao, "React: Synergizing reasoning and acting in language models," in *The eleventh international conference on learning representations*, 2022.

[12] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.

[13] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European control conference (ECC)*. Ieee, 2019, pp. 3420–3431.

[14] 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.

[15] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 322–328.

[16] S. Dean, A. Taylor, R. Cosner, B. Recht, and A. Ames, "Guaranteeing safety of learned perception modules via measurement-robust control barrier functions," in *Conference on Robot Learning*. PMLR, 2021, pp. 654–670.

[17] J. Breeden and D. Panagou, "Robust control barrier functions under high relative degree and input constraints for satellite trajectories," *Automatica*, vol. 155, p. 111109, 2023.

[18] E. Garone, S. Di Cairano, and I. Kolmanovsky, "Reference and command governors for systems with constraints: A survey on theory and applications," *Automatica*, vol. 75, pp. 306–328, 2017.

[19] E. G. Gilbert and I. Kolmanovsky, "Fast reference governors for systems with state and control constraints and disturbance inputs," *International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal*, vol. 9, no. 15, pp. 1117–1141, 1999.

[20] I. Filippidis and R. M. Murray, "Layering assume-guarantee contracts for hierarchical system design," *Proceedings of the IEEE*, vol. 106, no. 9, pp. 1616–1654, 2018.

[21] J. Panerati, H. Zheng, S. Zhou, J. Xu, A. Prorok, and A. P. Schoellig, "Learning to fly—a gym environment with pybullet physics for reinforcement learning of multi-agent quadcopter control," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 7512–7519.

[22] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016.