

Formation-Aware Dynamic Movement Primitives

A Master Thesis

M.Sc. Robotics

Alexander James Becoy

Delft University of Technology



Formation-Aware Dynamic Movement Primitives

A Master Thesis

by

Alexander James Becoy
(4904494)

Student number: 4904494
Thesis Committee: Dr. ir. L. Peternel, TU Delft, chair
Dr. R.T. Rajan, TU Delft, supervisor
Prof.dr.ir. M. Wisse, TU Delft
Project Duration: July, 2025 - June, 2026
Faculty: Mechanical Engineering, Delft
Electrical Engineering, Mathematics and Computer Science, Delft
Department: Cognitive Robotics
Signal Processing Systems, Microelectronics

Cover: TU Delft Moonshot Program



Preface

This project marks the culmination of my four-year journey toward becoming a robotics engineer at TU Delft. It has been a long and intense road (especially since I practically worked on two different theses) but I don't regret a single moment. The extended time allowed me to build an incredible network, find great career opportunities, and make lifelong friends. From building a community on the board of the Robotics Student Association and developing lower-limb exoskeletons with Project MARCH, to creating a humanoid robot hand for an art film and field-testing a quadruped in a Sicilian lava tube, the experiences have been endless. Along the way, I've gained technical expertise in robotics, space exploration, and wireless communication, while sharpening my research, problem-solving, and presentation skills. Most importantly, this journey helped me understand myself and what I truly want to achieve in life.

I am incredibly grateful to my supervisors, Luka Peternel and Raj Thilak Rajan, for your brilliance mentorship, endless patience, and unwavering encouragement throughout this challenging thesis journey. You did far more than just oversee my progress, you consistently pushed me to think deeper, provided profound technical insights, and trusted me with incredible opportunities through Moonshot and TRACI. Your ability to balance rigorous academic critique with genuine support was instrumental in helping me shape my raw ideas into an engineering project I am truly proud of.

I also owe a big thank you to Lunar Zebro and the TRACI team, Gijs Zijderveld, Robert Simonescu, and Aniruddh Pandey, for their assistance and camaraderie. It was an honor working together to get the rovers running. Watching them all work at the demonstration was a highlight, even if half of them predictably broke down by the end. Alas, the classic fate of robots at live demos. I look forward to seeing what both teams achieve next.

Special thanks to SpectX, Akhilesh Goveas, Mani Kamalzadeh, and the rest of the team. Thank you for providing exciting drone solutions that allowed me to apply my skills in the industry, while also helping me fund my living expenses.

I am equally grateful to my friends, Zhengyang Lu and Luis Roberto Alcantara Flores, for the fun and insightful times we shared while working on our theses and collaborating at SpectX. I wish you both the absolute best with your challenges.

Lastly, I want to thank my friends, family, and especially my girlfriend. Your unwavering love and support kept me going. Without you, this long journey would have been a much darker and more tedious ascent.

And to everyone else. Thank you.

*Alexander James Becoy
Delft, June 2026*

Contents

Preface	i
1 Introduction	1
2 Preliminaries	3
3 Methods	4
3.1 Trajectory Generation from Operator-Defined Via-Regions	4
3.2 Path Planning Based On Dynamic Movement Primitives	5
3.3 Multi-Agent Planning and External Forcing Terms	5
3.3.1 Formation Maintenance	5
3.3.2 Obstacle Avoidance	6
3.3.3 Inter-Agent Collision Avoidance	6
3.3.4 Connectivity Maintenance	7
3.4 Adaptive Temporal Scaling	7
3.4.1 Proximity-based Temporal Adaptation	7
3.4.2 Distributed Temporal Synchronization	7
3.4.3 Effects on DMP Dynamics	7
4 Experiments	8
4.1 Supervisory teleoperated control experiment	8
4.1.1 Experimental Setup	8
4.1.2 Tasks	9
4.1.3 Metrics	9
4.1.4 Results	10
4.2 Local autonomy in cluttered environments experiment	12
4.2.1 Experimental Setup	12
4.2.2 Tasks	12
4.2.3 Metrics	12
4.2.4 Results	14
5 Discussion	15
A Lunar Zebro Microrover Platform	17
B Ultra-Wideband Calibration	19
C Proof of Stability and Convergence	28
D Supplementary Results and Figures	30

Formation-Aware Dynamic Movement Primitives: Enabling Adaptive and Cooperative Leader-Follower-Based Multi-Robot Navigation For Shared Teleoperated Control

Alexander James Becoy^{1,2}

Supervised by:

Luka Peternel¹, Raj Thilak Rajan²

Abstract—Coordinated multi-robot networks are essential for planetary exploration. However, achieving fully autonomous operation in extreme extraterrestrial environments remains challenging due to stringent safety, stability, and reliability requirements. This paper presents a semi-autonomous supervisory teleoperation framework that combines high-level human oversight with decentralized multi-agent coordination. A remote operator specifies mission objectives and intermediate via-regions, while the robotic network autonomously performs real-time, collision-free navigation. To address GNSS-denied conditions, Ultra-Wideband (UWB) technology is employed for relative localization, requiring structured formation maintenance for hardware calibration. The proposed control architecture integrates Dynamic Movement Primitives (DMPs) for trajectory tracking with Affine Formation Control (AFC) for formation preservation, supplemented by adaptive collision avoidance mechanisms, and adaptive temporal scaling. Experimental validation demonstrates sub-meter tracking and formation accuracy, while extensive Monte Carlo simulations confirm robust destination reachability, formation maintenance, and network connectivity in cluttered environments.

I. INTRODUCTION

In the early stages of extraterrestrial exploration, robotic systems play a critical role in advancing scientific understanding and preparing infrastructure for future human missions, due to their scalability, redundancy, and distributed coordination capabilities [1]. These collectively enable robust infrastructure preparation [2], [3], scientific measurements collection [4], [5], and the identification of potential sites for human habitation [6], [7]. To fulfill these objectives, each robotic network must be capable of navigating unknown and inaccessible environments [8]. Although significant progress has been made in the research and development of such cooperative robotic networks, fully autonomous ones remain insufficiently mature for operation in extreme environments such as outer space [9], [10]. This limitation arises from stringent technical readiness requirements [11], particularly with respect to ensuring system stability, effective risk mitigation, and cost-effectiveness, which are critical considerations for the deployment and operation of space-borne technologies [12].

To mitigate unnecessary damage and unintended behaviors during navigation across the harsh extraterrestrial environments, incorporating remote human-in-the-loop within the

robotic networks enhances operational safety, particularly in long-distance operation [10], [13]. This approach capitalizes on human cognitive adaptability in complex and uncertain scenarios, enabling timely intervention during hazardous traversals or facilitating dynamic trajectory adjustments to explore regions of potential interest that autonomous systems may not independently identify. In practice, such oversight is commonly realized through supervisory teleoperation [14]–[16], wherein a human operator provides high-level commands, such as specifying sequence of interesting destinations for the robots to visit and regions to avoid, while also reducing operator effort and cognitive workload. Simultaneously, the robotic network maintains low-level autonomy, allowing it to navigate challenging terrain in real time, perform collision avoidance against unmodeled obstacles, and interact with other agents, as illustrated in Fig. 1. The robotic network further transmits sparse telemetry and status information to the operator, thereby providing the telepresence necessary to maintain situational awareness of the system and its surrounding environment. Consequently, this framework establishes a semi-autonomous operational architecture that effectively combines human strategic oversight with autonomous robotic execution.

To enable effective real-time autonomous navigation in multi-robot systems operating in inaccessible, unknown, and GNSS-denied environments, in addition to enforcing goal-directed motion toward target positions while ensuring collision avoidance, the incorporation of wireless inter-agent communication within the local network produces emergent cooperative behaviors that can significantly enhance navigational performance [17]–[20]. These include the exchange of situational awareness regarding both the agents' dynamical states and their environment [21], as well as the use of relative localization techniques based on affine formation structures [22], [23]. However, several challenges remain in achieving robust communication, particularly with respect to timing between distributed devices, and coping with antenna-induced delays that affect the time-of-flight (ToF) measurements. Reducing these low-level uncertainties necessitates calibration among devices, which can be achieved by maintaining formations with known inter-agent distances [24], [25]. Nevertheless, preserving rigid formations in extremely cluttered environment introduces an additional layer of difficulty, as it constrains both motion planning and coordination under environmental and communication limitations.

Real-time coordinated multi-robot navigation has been extensively investigated across literature, each presenting distinct

¹Department of Cognitive Robotics, Faculty of Mechanical Engineering, Delft University of Technology, The Netherlands.

²Signal Processing Systems Research Group, Department of Microelectronics, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, The Netherlands.




Fig. 1: System pipeline of a human operator supervising a remote robotic network through a long-range space communication link via a lunar lander. The operator issues high-level commands, such as a target destination, while the robotic network autonomously executes collision-avoiding navigation from position *A* to position *C* through an obstacle-rich region in region *B*, as traced by the blue trajectory. Sparse telemetry and status updates are returned asynchronously across the network to maintain operator situational awareness, while the orange lines denote undirected communication links maintained between neighboring agent pairs.

trade-offs between computational overhead, safety guarantees, structural flexibility, and model dependency. Centralized methods, such as Artificial Potential Fields (APF) applied to human-directed multi-drone systems [26], offer high run-time efficiency but suffer from a single point of failure, constrained scalability, and a heavy reliance on continuous human intervention and design. To eliminate centralized bottlenecks, distributed optimization-based frameworks have been proposed. Distributed Model Predictive Control (D-MPC) formulations [17] enable independent, constraint-aware trajectory planning over finite horizons, but introduce substantial online computational overhead due to repeated optimization and are highly sensitive to accurate system models and optimization infeasibility. Conversely, distributed Control Barrier Functions (CBFs) unified with Control Lyapunov Functions (CLFs) [18] leverage forward-invariant safe sets to provide mathematically rigorous collision avoidance. However, standard CBF architectures remain model-dependent, can restrict solution optimality by overly constraining the control space, and demand meticulous tuning to avoid deadlocks.

To bypass explicit system modeling while retaining fast online inference, learning-based methodologies and stable dynamical system behaviors have emerged. Multi-Agent Reinforcement Learning (MARL) policies optimized via Proximal Policy Optimization (PPO) [19] can implicitly encode complex formation and avoidance objectives through reward functions. Nonetheless, MARL suffers from high policy variance, heavy offline training requirements, and a lack of formal safety guarantees under environmental distribution shifts. As a non-learning alternative that requires no explicit plant modeling,

Dynamical Movement Primitives (DMPs) [27] structure trajectories as stable dynamical systems with guaranteed destination convergence, augmented by analytical forcing terms for multi-manipulator synchronization [28] and rover navigation [29]. Yet, current DMP architectures are largely bounded by rigid formation tracking or point-mass obstacle potentials that fail to capture volumetric shapes and dynamics, restricting fluid formation deformation and scaling around dense obstacles. Essentially, there remains a notable gap in literature for a scalable, distributed DMP framework that seamlessly unifies complex volumetric obstacle avoidance with resilient, shape-adaptive affine formation recovery in highly cluttered, safety-critical environments.

Motivated by the need to incorporate human supervision, eliminate single points of failure, reduce reliance on explicit system modeling, ensure both safety and guaranteed convergence towards target states while maintaining a desired formation, we propose a novel framework for coordinated multi-agent navigation in challenging environments such as those encountered on the lunar surface. The proposed approach integrates shared teleoperation with DMPs, volumetric obstacle avoidance that explicitly captures obstacle geometry and dynamics, and affine formation structures to enable flexible transformation of formation shapes, thereby achieving robust, scalable, and adaptive coordination under unknown and constrained settings. The main contributions of this work include: 1) a DMP-based formation control method in which leader agents generate motion primitives that are propagated to follower agents through geometric transformations, which enables coordinated navigation while preserving formation struc-

TABLE I: Comparative Analysis of Multi-Agent Autonomous Navigation Methodologies

Methodology	Model Dependency	Computational Cost	Safety Attributes	Formation Flexibility
Centralized APF [26]	Low	Low	Local Minima Prone	Rigid, Manually Designed
Distributed MPC [17]	High	High	Feasibility Dependent	Constrained Bounded
Distributed CBF-CLF [18]	High	Moderate	Mathematically Rigorous	Conservatively Limited
MARL / PPO [19]	Low	Low (Online Inference)	Empirical / Weak	Policy-Dependent
Standard DMP [27], [29]	Low	Low	Asymptotically Stable	Rigid / Non-Adaptive
Proposed Framework	Low	Low	Adaptive Potential Guarded	Highly Adaptive (Affine)

ture under dynamic conditions; 2) a shared-control paradigm that incorporates high-level human guidance through intermediary regions while maintaining autonomous low-level control for navigation, formation maintenance, and safety, thereby balancing operator workload with system robustness; and 3) comprehensive validation through extensive simulations in both obstacle-free and highly cluttered environments, with performance evaluated using metrics such as reachability, formation accuracy, safety margin, and computational efficiency.

The remainder of the paper is structured as follows, Section B.2 presents the problem statement and relevant definitions. Section III introduces the proposed extension of DMPs for coordinated multi-agent navigation. Section IV details the experimental evaluation, and Section V concludes with a discussion of the results and implications.

Notation: Scalars are denoted by standard letters (e.g., x or X), vectors by bold lowercase letters (e.g., $\mathbf{x}$), and matrices by bold uppercase letters (e.g., $\mathbf{X}$). Sets are represented by calligraphic symbols (e.g., $\mathcal{X}$). The Kronecker product, element-wise division, and transpose operators are denoted by $\otimes$, $\oslash$, and $(\cdot)^\top$, respectively. Discrete-time quantities are indexed by superscripts $(\cdot)^{(k)}$, where k denotes the time step. Agent- and obstacle-specific variables are distinguished by subscripts (e.g., $\mathbf{x}_i$), whereas module-specific quantities are indicated by textual superscripts (e.g., $\mathbf{x}^{\text{module}}$). The $N \times N$ identity matrix is denoted by $\mathbf{I}_N$, and the zero matrix of dimension $A \times B$ is denoted by $\mathbf{0}_{A \times B}$.

II. PRELIMINARIES

We consider a semi-autonomous multi-agent robotic system operating on the lunar surface under remote human supervision, as illustrated in Fig. 1. The operating environment is highly unstructured, unknown, and inaccessible, with terrain characterized by natural hazards, such as rocks and craters [30]. To narrow the scope of this work to multi-agent navigation and coordination, we assume the availability of a GNSS-like localization infrastructure, such as orbital satellites, which provides the position of each agent, denoted by $\mathbf{z}$, relative to a global inertial reference frame, for instance, with the lunar lander coordinate frame serving as the origin.

The robotic network is modeled as a multi-agent system composed of N agents operating in a D -dimensional Euclidean space. Each rover is modeled as a circular rigid body with radius r^r and is equipped with onboard sensing and communication capabilities. Obstacles can be detected within a sensing radius r^s , while inter-agent communication is supported within a communication radius r^c , thereby defining the local perception and interaction boundaries of each agent

and governing information exchange within the network. The rover specifications are further described in Appendix A.

The interaction structure among the agents is represented by an undirected communication graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ denotes the set of agents with cardinality $|\mathcal{V}| = N$, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ denotes the set of communication links between pairs of agents, characterizing the information exchange. For each agent $i \in \mathcal{V}$, the corresponding neighbor set is defined as $\mathcal{N}_i = \{j \in \mathcal{V} \setminus i : (i, j) \in \mathcal{E}\}$.

To ensure synchronized operation, particularly during hardware calibration procedures such as ultra-wideband (UWB) calibration, the robotic network is required to achieve a predefined geometric configuration, preferably before and after navigation tasks, as detailed in Appendix B. A formation is formally defined as $\mathcal{F} = (\mathcal{G}, \mathbf{P})$, where $\mathcal{G}$ denotes the communication topology of the network and $\mathbf{P} = [\mathbf{p}_1^\top, \mathbf{p}_2^\top, \dots, \mathbf{p}_N^\top] \in \mathbb{R}^{D \times N}$ represents the configuration of the N agents [31]. The nominal formation is denoted by $\mathcal{F}^* = (\mathcal{G}, \mathbf{P}^*)$, where $\mathbf{P}^*$ specifies the target geometric configuration.

To facilitate affine formation structures, affine formation control is incorporated which typically imposes a hierarchical leader-follower architecture on the network [22]. The set of leaders is denoted by $\mathcal{V}_L$, with cardinality $|\mathcal{V}_L| = N_L = D+1$, satisfying the geometric requirements for AFC. The remaining agents form the set of followers, defined as $\mathcal{V}_F = \mathcal{V} \setminus \mathcal{V}_L$, with cardinality $|\mathcal{V}_F| = N_F$. This hierarchical organization enables efficient distributed coordination while preserving the geometric properties of the formation throughout the navigation process. Since the robotic network is modeled as an undirected graph, implying bidirectional communication between connected agents, realizability and stabilizability of a distance-constrained formation are guaranteed if and only if the underlying framework is universally rigid. This means that any distance-preserving realization of the formation in any Euclidean space is unique up to rigid transformations (translation, rotation, and shearing), and therefore coincides with the original formation.

To generate task-level motion, leader agents are equipped with DMPs, which provide a high-level trajectory generation framework for goal-directed navigation and demonstration imitation. The DMP framework is further augmented with APF components to incorporate repulsive behaviors for obstacle avoidance and mitigation of inter-agent collisions [27]. Furthermore, it is assumed that both the initial and final positions of the generated trajectories are free from disturbances and environmental hazards, thereby ensuring reliable convergence of the agents toward their designated configurations. Follower agents are likewise equipped with DMP-based guidance terms that steer them toward their designated target positions. How-

ever, unlike the leaders, followers do not directly imitate the demonstrated trajectories. Instead, the relative contributions of DMP and AFC objectives are weighted such that formation maintenance remains the dominant objective, particularly during the final stages of navigation.

Consequently, the DMP component in the follower agent serves primarily as a secondary guidance mechanism, while the AFC term ensures accurate convergence to the prescribed geometric configuration. Since the leaders define the global motion through their reference trajectories and induce the corresponding geometric transformation of the nominal formation $\mathcal{F}^*$, followers are only required to track the evolving formation rather than replicate the demonstration itself, thereby simplifying their role while preserving coordinated group behavior. Both leader and follower agents are equipped with APF-based collision avoidance mechanisms to ensure safe operation in cluttered environments. Although these avoidance behaviors may temporarily perturb the formation structure, the proposed control framework guarantees recovery of the desired formation once the disturbances are removed. Since strict formation maintenance is primarily required outside active navigation phases, such as during calibration or stabilization when the rovers are stationary, such transient deviations are considered acceptable.

III. METHODS

The proposed framework, illustrated in Fig. 2, introduces a distributed control architecture for semi-autonomous multi-agent navigation and formation maintenance in unstructured lunar environments. At a high-level, the method integrates (i) operator-defined task specification through intermediary regions in the teleoperation interface, (ii) DMP-based path planning for leader agents, (iii) multi-agent coordination with external forcing terms for collision avoidance and formation maintenance, and network coordination through (iv) adaptive temporal scaling for risk-aware motion modulation. Together, these components form a unified pipeline that maps high-level human intent into safe, coordinated, and geometrically consistent multi-robot motion. The stability and convergence of the proposed framework are proven in Appendix C.

A. Trajectory Generation from Operator-Defined Via-Regions

Prior to deployment, a human operator specifies a mission plan consisting of an ordered sequence of intermediary regions, which we denote as *via-regions*, through which the multi-agent system must pass en route to its final destination, together with a nominal formation $\mathcal{F}^*$ that defines the desired spatial arrangement and communication between the agents. This operator-provided information serves as the high-level demonstration from which the controllers of the leaders learn their reference behavior, following the learning from demonstration (LfD) paradigm [32], [33].

Consider the first leader, L_1 , defines the reference frame of formation $\mathcal{F}^*$. Given consecutive via-region centers $\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{N_W} \in \mathbb{R}^D$, where $\mathbf{w}_0$ is the starting position of the formation frame with respect to L_1 's reference frame, and $\mathbf{w}_{N_W}$ is its final position. A set of reference trajectories




Fig. 2: Control block scheme diagram depicting the proposed framework. The human operator (purple) provides via-regions that the robotic network must visit to the teleoperation interface (green) which generates the corresponding reference trajectories for the leader agents. Blue boxes represent the agents' low-level modules such as sensing, locomotion, and communication capabilities. Yellow and red boxes correspond to the main contribution, where yellow describes methods related to path planning per agent, and red enables distributed coordination between agents.

$\mathcal{R} = \{\mathbf{R}^{(1)}, \mathbf{R}^{(2)}, \dots, \mathbf{R}^{(N_R)}\}$ is generated for each leader, where $\mathbf{R}^{(\rho)} \in \mathbb{R}^{D \times N_L}$, $\forall \rho = 1, 2, \dots, N_R$, collects the reference setpoints of all N_L leaders at sample ρ , such that

$$\mathbf{R}^{(r)} = \begin{bmatrix} \mathbf{r}_1^{(\rho)\top} & \mathbf{r}_2^{(\rho)\top} & \dots & \mathbf{r}_{N_L}^{(\rho)\top} \end{bmatrix}^\top \in \mathbb{R}^{D \times N_L}, \quad (1)$$

where N_R denotes the total number of trajectory setpoints.

The trajectory in the formation's reference frame is constructed as a piecewise cubic Hermite spline passing through the prescribed waypoint sequence [34]. Arc-length parameterization is employed so that the parameter $\xi \in [0, 1]$ advances proportionally to the Euclidean distance along the path, ensuring approximately uniform speed across segments of unequal length. Tangent directions at interior waypoints are assigned using the Catmull-Rom rule, whereby each tangent is the normalized average of the incoming and outgoing segment directions, guaranteeing C^1 continuity of position and first derivative throughout the path. The trajectories of the remaining leaders $L_2, \dots, L_{N_L}$ are then derived analytically from L_1 's path and according to their respective offsets in $\mathbf{P}^*$. The instantaneous heading of the group's trajectory $\psi(t) \in S^1$ is extracted from the analytic spline derivative and smoothed with a Gaussian kernel to ensure temporal continuity in the heading profile, after which each leader's world frame position is computed as

$$\mathbf{r}_i(t) = \mathbf{r}_{L_1}(t) + R(\psi_{\text{eff}}(t))\mathbf{t}_i(t) \in \mathbb{R}^D, \quad (2)$$

where t is time, $R(\cdot) \in \mathbb{R}^{D \times D}$ is the rotation matrix, $\mathbf{t}_i(t)$ is the body-frame offset of leader i from L_1 interpolated linearly from its initial to its final value, and $\psi_{\text{eff}}(t)$ is a raised-cosine fade of $\psi(t)$ that vanishes at both endpoints to guarantee exact position boundary conditions without velocity discontinuities. T is the total time required to complete the demonstration.

The resulting set of reference trajectories $\mathcal{R}$ are then presented to the DMPs of the respective leaders as demonstrations. Each DMP fits a set of forcing functions weights $\zeta \in \mathbb{R}$ via least-square regression, such that integrating the DMP from the prescribed initial state reproduces the demonstrated path. The learned weights encode the shape of the reference trajectory in a time-scaled, coordinate-invariant form, enabling the DMP to generalize to changes in goal position and execution speed [27].

While the learned forcing term provide leaders with globally consistent reference path, they are insufficient to handle small-scale terrain features, such as rocks and small craters, that are not visible from orbital survey imagery and therefore absent from the operator-provided mission plan. To address this, supplementary forcing terms are incorporated into the DMP integration runtime, as similarly presented in [35], [36]. These terms generate repulsive accelerations based on proximity to detected obstacles, allowing each leader to deviate locally from the reference trajectory as necessary while continuing to track its overall goal. This combination of offline demonstration learning and online reactive forcing thus enables the leaders to maintain global path coherence while autonomously negotiating unforeseen impediments in the terrain.

B. Path Planning Based On Dynamic Movement Primitives

To generate smooth, goal-directed trajectories for the multi-agent system, each leader agent is equipped with a discrete DMP. Discrete DMP encodes point-to-point motions, such as reference trajectories obtained from demonstrations, into a stable nonlinear dynamical systems [37]. This representation ensures robustness to perturbations while allowing adaptation to new and goal configurations.

Let $\mathbf{z}_i, \dot{\mathbf{z}}_i, \ddot{\mathbf{z}}_i \in \mathbb{R}^D$ denote the position, velocity, and acceleration of leader agent i , respectively. The DMP formulation consists of a transformation system (3) and (4), governing the motion dynamics, and a canonical system (5), which provides a temporal parameterization. The overall system is defined by a set of ordinary differential equations (ODEs):

$$\tau_i \dot{\mathbf{v}}_i = \mathbf{A}_z (\mathbf{B}_z (\mathbf{z}_i^{(T)} - \mathbf{z}_i) - \mathbf{v}_i) + \mathbf{f}_i(s_i) + \Phi_i^{\text{form}} + \Phi_i^{\text{obs}} + \Phi_i^{\text{agent}} + \Phi_i^{\text{conn}}, \quad (3)$$

$$\tau_i \dot{\mathbf{z}}_i = \mathbf{v}_i, \quad (4)$$

$$\tau_i \dot{s}_i = -\alpha_s s_i, \quad (5)$$

where the matrices $\mathbf{A}_z, \mathbf{B}_z \in \mathbb{R}^{D \times D}$ are positive semi-definite and determine the behavior of the underlying second-order system. In the absence of forcing and coupling terms, i.e., when $\mathbf{f} + \Phi \equiv \mathbf{0}$, f , critical damping is imposed, $\mathbf{B}_z = \sqrt{2}\mathbf{A}_z$, to ensure convergence to the equilibrium $(\mathbf{z}_i, \dot{\mathbf{z}}_i) = (\mathbf{z}_i^{(T)}, \mathbf{0})$. The scalar $\tau_i \in \mathbb{R}_{++}$ provides temporal scaling, allowing modulation of the execution speed.

The phase variable $s_i \in (0, 1]$ acts a surrogate for time $t \in [0, T]$, where T denotes the total duration of the demonstrated trajectory. It evolves monotonically from 1 to 0. This decay ensures that the nonlinear spring forcing term vanishes near the goal, yielding smooth convergence with zero terminal velocity. $\Phi_i^{\text{form}}, \Phi_i^{\text{obs}}, \Phi_i^{\text{agent}}, \Phi_i^{\text{conn}} \in \mathbb{R}^D$ represent external influences corresponding to formation maintenance, obstacle avoidance, inter-agent collision avoidance, and connectivity maintenance, respectively.

The nonlinear forcing term $\mathbf{f}_i$ encodes the shape of the desired trajectory and is learned from demonstrations. Each component $f_i^b \in \mathbf{f}_i \forall b \in [1, N_B]$ is represented as a weighted combination of basis functions:

$$f_i^b(s_i) = \frac{\sum_{b=0}^{N_B} {}^p\zeta_b \sigma_b(s_i)}{\sum_{b=0}^{N_B} \sigma_b(s_i)} s_i \quad (6)$$

where ${}^p\zeta_b \in \mathbb{R}$ are the weights, p is a component in the D -dimensional space, and $\sigma_b(s_i)$ is a Gaussian radial basis function defined as

$$\sigma_b(s_i) = \exp(-h_b(s_i - c_b)^2). \quad (7)$$

The centers, c_b , and widths, h_b , of the basis functions are chosen as

$$c_b = \exp\left(-\alpha_s b \frac{T}{N_B}\right), \quad b = 0, 1, \dots, N_B \quad (8)$$

$$h_b = \begin{cases} \frac{1}{(c_{b+1} - c_b)^2} & b = 0, 1, \dots, N_B - 1, \\ h_{N_B-1} & b = N_B. \end{cases} \quad (9)$$

DMP operates in two phases. During the *learning phase*, the weights ${}^p\zeta_b$ are computed from a set of reference trajectories $\mathcal{R}$ using linear regression, such that the forcing term accurately reproduces the demonstrated motion. During the *execution phase*, the learned dynamical system is integrated from the given initial and goal states, producing trajectories that preserve the qualitative shape of the demonstrations while adapting to new boundary conditions. This property enables flexible and robust trajectory generation for leaders in dynamic and uncertain environments. Formally, for each agent i at time step k , solving (3)-(5) yields the next desired state, namely the position and velocity pair $(\hat{\mathbf{z}}_i^{(k+1)}, \dot{\hat{\mathbf{z}}}_i^{(k+1)})$.

C. Multi-Agent Planning and External Forcing Terms

1) *Affine Formation Control*: According to rigidity theory, a multi-agent formation can preserve its geometric structure through a set of inter-agent constraints that uniquely characterize the desired formation. In this context, the stress matrix $\Omega \in \mathbb{R}^{N \times N}$, with entries $\omega_{i,j}$ encodes the intrinsic geometric relationships of the nominal universally-rigid formation $\mathcal{F}^* = (\mathcal{G}, \mathbf{P}^*)$. These coefficients impose equilibrium constraints that lie in the nullspace of the rigidity matrix

$$\Omega \mathbf{P}^{*\top} = \mathbf{0}_{N \times D}, \quad (10)$$

ensuring consistency of the formation structure. As a consequence, Ω guarantees invariance of $\mathcal{F}^*$ under allowable affine

transformations, preserving the formation up to translation, rotation, scaling, and shearing.

Formation maintenance within the multi-agent system is achieved via an AFC law for double-integrator agents with time-varying velocities. At discrete time instant k , the formation component of the control input for follower agent i is given by [38]:

$$\begin{aligned} \Phi_i^{\text{form},k} = & -\frac{1}{\gamma_i} \sum_{j \in \mathcal{N}_i} \omega_{i,j} (\mathbf{A}_i^{\text{form}}(\mathbf{z}_i^{(k)} - \mathbf{z}_j^{(k)}) \\ & + \mathbf{B}_i^{\text{form}}(\dot{\mathbf{z}}_i^{(k)} - \dot{\mathbf{z}}_j^{(k)}) - \ddot{\mathbf{z}}_j^{(k)}), \end{aligned} \quad (11)$$

where $\gamma_i = \sum_{j \in \mathcal{N}_i} \omega_{i,j} \in \mathbb{R}_{++}$, and $\mathbf{A}_i^{\text{form}}, \mathbf{B}_i^{\text{form}} \in \mathbb{R}^{D \times D}$ are positive semi-definite gain matrices governing position- and velocity-level coupling.

Although this framework guarantees convergence to the affine shape $\mathbf{P}^*$ under nominal conditions [22], the formation is not strictly enforced during all phases of operation. In particular, since coordinated formation is primarily required outside active navigation phases (e.g., during calibration or structured deployment stages), temporary deviations from $\mathbf{P}^*$ are permitted in the presence of disturbances. Once such disturbance vanishes, the control law ensures asymptotic recovery of the affine shape $\mathbf{P}^*$.

2) *Obstacle Avoidance*: The lunar environment contains static obstacles, primarily rocks and craters, which are assumed to remain fixed in the world frame. To ensure safe navigation, each agent incorporates a reactive obstacle avoidance mechanism based on dynamic potential fields.

Let $\mathbf{z}_m \in \mathbb{R}^D$ denote the position of obstacle m . Each obstacle is modeled as a repulsive *dynamic volumetric potential field*, which generates avoidance forces that depend on the relative position and velocity of the interacting agent, as well as the obstacle's geometric volume [39]. Such a formulation activates repulsive effects primarily when the agent is moving toward the obstacle, thereby reducing unnecessary trajectory deviations and oscillatory behavior. The resulting repulsive force acting on agent i due to obstacle m is given by

$$\begin{aligned} \varphi_{i,m}(\mathbf{z}_i, \dot{\mathbf{z}}_i) = & -\frac{\lambda \|\dot{\mathbf{z}}_i\| (-\cos \theta_{i,m})^{\beta-1}}{C_m^\eta(\mathbf{z}_i)} \times \\ & \left(-\beta \nabla_{\mathbf{z}_i} (\cos \theta_{i,m}) + \frac{\eta \cos \theta_{i,m}}{C_m(\mathbf{z}_i)} \nabla_{\mathbf{z}_i} C_m(\mathbf{z}_i) \right), \end{aligned} \quad (12)$$

where $\lambda > 0$, $\beta \geq 0$, and $\eta > 0$ are design parameters that regulate the magnitude and spatial influence of the repulsive field. The function $C_m(\cdot) \in \mathbb{R}_+$ denotes the isopotential function associated with obstacle m , which ensures that the repulsive effect decays with increasing distance between agent i and the obstacle. Furthermore, the magnitude of the repulsive field scales with the agent's speed $\|\dot{\mathbf{z}}_i\|_2$ and vanishes when the agent is stationary.

For an ellipsoidal obstacle, the isopotential function with respect to agent i is defined as

$$C_m(\mathbf{z}_i) = \|(\mathbf{z}_i - \mathbf{z}_m) \odot \boldsymbol{\ell}_{m,\text{eff}}\|_2^2, \quad (13)$$

where $\mathbf{z}_m$ denotes the centroid of obstacle m , and $\boldsymbol{\ell}_{m,\text{eff}}$ represents the effective semi-axis lengths of the corresponding ellipsoid, computed by inflating the obstacle's semi-axis lengths $\boldsymbol{\ell}_m$ by the agent's physical radius r^r and a minimum clearance parameter $c^{\text{obs}} \geq 0$.

The term $\theta_{i,m}$ denotes the angle between the agent velocity $\dot{\mathbf{z}}_i$ and the relative position vector $\mathbf{z}_i - \mathbf{z}_m$, defined as

$$\cos \theta = \frac{\langle \nabla_{\mathbf{z}} C_m(\mathbf{z}_i), \dot{\mathbf{z}}_i \rangle}{\|\nabla_{\mathbf{z}} C_m(\mathbf{z}_i)\| \|\dot{\mathbf{z}}_i\|}. \quad (14)$$

This directional term ensures that the repulsive force is significant only when the agent is moving toward the obstacle (i.e., $\cos \theta_{i,m} < 0$, and negligible otherwise).

The function $C_m(\mathbf{z}_i)$ represents an implicit superquadric distance function describing the geometry of obstacle m . Its gradient $\nabla_{\mathbf{z}} C_m(\mathbf{z})$ provides the outward normal direction of the obstacle surface, enabling anisotropic and geometry-aware avoidance. The exponent η governs how sharply the repulsive force increases as the agent approaches the obstacle boundary.

At each discrete timestep k , agent i considers only obstacles within its sensing radius r^s . This reduces the number of obstacles needed to be computed, therefore improving the computational efficiency. The corresponding set is defined as

$$\mathcal{O}_i^{(k)} = \{\forall m \in \mathcal{O} \mid \|\mathbf{z}_m - \mathbf{z}_i\|_2 \leq r^s\} \quad (15)$$

where $\mathcal{O}$ denotes the set of all obstacles in the environment.

The overall obstacle avoidance forcing term acting on agent i is obtained by summing the contributions from all sensed obstacles:

$$\Phi_i^{\text{obs},(k)} = \sum_{m \in \mathcal{O}_i^{(k)}} \varphi_{i,m}^{(k)}(\mathbf{z}_i - \mathbf{z}_m, \dot{\mathbf{z}}_i). \quad (16)$$

This velocity-aware and geometry-consistent formulation enables smooth and reactive obstacle avoidance, making it well-suited for integration within the DMP-based control framework while preserving trajectory coherence and safety.

3) *Inter-Agent Collision Avoidance*: In addition to obstacle avoidance, safe operation of the multi-agent system requires explicit consideration of inter-agent interactions, as agents exhibit relative motion with respect to one another. Consequently, collision avoidance must account not only for relative positions but also for relative velocities to ensure dynamically consistent safety behavior.

Since agents are equipped with wireless communication capabilities, the need for explicit sensing of neighboring agents is eliminated. Instead, each agent can directly access the state information (positions and velocities) of agents within its communication neighborhood. This communication-driven paradigm is particularly advantageous in formation-based control settings, where agents are already structured according to a predefined interaction topology. As a result, collision avoidance can be restricted to non-adjacent agents while preserving formation constraints.

Each neighbors j is modeled as a repulsive dynamic volumetric potential field, extending the formulation used for static obstacles to moving agents. The corresponding repulsive forcing term acting on agent i due to agent j is defined as

$$\varphi_{i,j}(\mathbf{z}_i - \mathbf{z}_j, \dot{\mathbf{z}}_i - \dot{\mathbf{z}}_j), \quad (17)$$

where the interaction depends on both the relative position and relative velocity. This velocity-dependent formulation enables prediction-aware repulsion, allowing agents to react not only to spatial proximity but also to potential future collisions. The effective semi-axis length of the corresponding isopotential function is given by $2r^r + c^{\text{agent}}$, where $c^{\text{agent}} \geq 0$ denotes a design parameter that accounts for the agent's safety margin.

At each discrete timestep k , agent i computes its inter-agent collision avoidance term by aggregating the contributions from all neighboring agents within its communication set $\mathcal{N}_i$:

$$\Phi_i^{\text{agent},(k)} = \sum_{j \in \mathcal{N}_i} \varphi_{i,j}^{(k)}(\mathbf{z}_i - \mathbf{z}_j, \dot{\mathbf{z}}_i - \dot{\mathbf{z}}_j). \quad (18)$$

This formulation ensures that inter-agent safety is maintained in a distributed manner while preserving formation coherence, as interactions are computed locally using communicated state information rather than explicit sensing. The resulting repulsive behavior is smooth, symmetric, and consistent with the dynamic volumetric potential field framework used for obstacle avoidance.

4) *Connectivity Maintenance*: When leader agents execute aggressive maneuvers to bypass expansive environmental hazards, they are often forced to displace significantly from one another. Because the localized target states of the follower agents scale proportionally relative to these leaders, this geometric expansion generates a larger macro-level affine shape footprint across the entire formation. As a consequence of this structural expansion, the physical distance between paired agents increases. If these localized distances transcend the maximum reliable communication range, denoted by r^c , critical topological links will break. This can drive the multi-agent system into a fragmented state, ultimately precipitating a total formation collapse.

To preserve network connectivity under large affine scaling conditions, a dedicated decentralized connectivity maintenance term, $\Phi_{i,j}^{\text{conn},(k)}$, is seamlessly integrated into the transformation system equations of the DMP. The connectivity restoration force acting on agent i with respect to its neighbor j is mathematically formalized as:

$$\Phi_{i,j}^{\text{conn},(k)} = \begin{cases} -\frac{1}{\gamma_i} \mathbf{A}^{\text{conn}}(\mathbf{z}_i - \mathbf{z}_j), & d_{i,j} \geq \alpha^{\text{conn}} r^c, \\ 0, & d_{i,j} < \alpha^{\text{conn}} r^c, \end{cases} \quad (19)$$

where $\mathbf{A}^{\text{conn}} \in \mathbb{R}^{D \times D}$ denotes a positive-definite control gain matrix governing the magnitude of the attractive restoration force. The scalar parameter $d_{i,j} = \|\mathbf{p}_i - \mathbf{p}_j\|_2$ represents the Euclidean distance between the physical positions of the agent pair (i, j) . The design parameter $\alpha^{\text{conn}} \in (0, 1)$ establishes a conservative safety threshold that dictates the activation onset of the connectivity mechanism prior to violating the physical communication boundary r^c .

When the localized multi-agent distance resides within the nominal safe envelope ($d_{i,j} \leq \alpha^{\text{conn}} r^c$), the connectivity term remains completely dormant ($\Phi_{i,j}^{\text{conn},(k)} = 0$), thereby allowing

the formation to deform fluidly according to the pure affine commands dictated by the leader trajectory. However, the instant an expansive maneuver stretches adjacent agents past the safety boundary ($d_{i,j} \geq \alpha^{\text{conn}} r^c$), the consensus-like attraction profile in (19) activates, dynamically constraining the agent trajectories to remain within their bounded communication envelope.

D. Adaptive Temporal Scaling

The interaction between DMP, formation control, and repulsive potential fields introduces strong nonlinear coupling in the closed-loop system. In particular, when agents are driven by high accelerations from (3), velocity-dependent repulsive forces associated with obstacles and inter-agent avoidance may increase sharply. This can lead to oscillatory behavior, including overshooting and rebound effects near obstacles and other agents, especially in dense or constrained environments.

To address this issue, we introduce an adaptive temporal scaling mechanism that regulates the execution speed of the entire dynamical system. The key idea is to slow down the evolution of both the transformation and canonical systems when agents approach critical regions, thereby improving stability and ensuring smoother interaction between attractive and repulsive terms. Recall that the DMP dynamics is given by (3)-(5), where the scalar τ_i governs the temporal evolution of both the transformation and canonical systems. Increasing τ_i slows down the system evolution, effectively stretching the time horizon of the execution phase, while preserving the spatial trajectory structure.

1) *Proximity-based Temporal Adaptation*: To dynamically modulate the formation's progression speed in response to environmental and structural constraints, an adaptive temporal scaling framework is established. The localized time-scaling factor is constructed as a function of multiple safety-critical distance metrics, capturing both obstacle proximity and inter-agent spatial separation.

Let $d_{i,j} = \|\mathbf{z}_i - \mathbf{z}_j\|_2 \in \mathbb{R}_+$ denote the Euclidean distance between agent i and an incident entity j , where j represents either an obstacle boundary or another agent. For connectivity maintenance, we define $d_{i,j} = \max\{0, \|\mathbf{z}_i - \mathbf{z}_j\|_2 - \alpha^{\text{conn}} r^c\}$, while for formation maintenance we define $d_{i,i^*} = \|\mathbf{z}_i - \mathbf{z}_i^*\|_2$, where $\mathbf{z}_i^*$ denotes the desired formation position of agent i . For notational convenience, let $d \in \{d_{i,j}, d_{i,i^*}\}$. The adaptive temporal scaling law for agent i is defined as

$$\tau_i^{\text{adapt}}(d) = \begin{cases} \tau_i^{\text{base}} + \tau_{\text{extra}} \left(1 - \frac{d}{d_{\text{slow}}}\right)^2, & 0 \leq d \leq d_{\text{slow}}, \\ \tau_i^{\text{base}}, & d > d_{\text{slow}}, \end{cases} \quad (20)$$

where $\tau_i^{\text{base}} \in \mathbb{R}_+$ denotes the nominal time-scaling factor governing baseline mission execution, $\tau_{\text{extra}} \in \mathbb{R}_+$ specifies the maximum admissible temporal augmentation (i.e., bounded deceleration), and $d_{\text{slow}} \in \mathbb{R}_+$ is the activation threshold distance beyond which no temporal modulation is applied.

Each agent computes multiple temporal scaling contributions:

$$\tau_i^{\text{obs}}, \quad \tau_i^{\text{agent}}, \quad \tau_i^{\text{conn}}, \quad \tau_i^{\text{form}}, \quad (21)$$

corresponding respectively to obstacle proximity, and inter-agent proximity, communication constraints, and formation tracking error. The final local temporal scaling is obtained via a conservative aggregation:

$$\tau_i = \max(\tau_i^{\text{obs}}, \tau_i^{\text{agent}}, \tau_i^{\text{conn}}, \tau_i^{\text{form}}). \quad (22)$$

2) *Distributed Temporal Synchronization*: To ensure consistent execution across the multi-agent system, the temporal scaling values are propagated over the communication graph. Each agent updates its value according to a max-consensus rule:

$$\tau_i^{(k+1)} = \max_{j \in \mathcal{N}_i \cup \{i\}} \tau_j^{(k)}. \quad (23)$$

This update follows the class of distributed extremum consensus algorithms, where agents asymptotically agree on the global maximum value through local exchanges [40]. Convergence of such max-consensus protocols over connected graphs has been extensively studied and shown to be achieved in a finite number of iterations bounded by the diameter of the communication graph [41]. Consequently, all agents eventually adopt the most conservative (slowest) temporal scale within the network.

3) *Effect on DMP dynamics*: The adaptive temporal scaling τ_i directly modulates both subsystems of the formulation:

- In the transformation system, τ_i scales the acceleration and velocity evolution, thereby reducing the magnitude of velocity-dependent interaction forces when safety constraints become active.
- In the canonical system, τ_i slows down the decay of the phase variable s_i , effectively delaying the attenuation of the forcing term $\mathbf{f}_i(s_i)$ and allowing smoother convergence under interaction constraints.

Consequently, the proposed adaptive mechanism dynamically regulates execution speed in response to environmental and interaction complexity, ensuring stable integration of DMP-driven motion generation, formation maintenance, and collision avoidance within a unified control framework.

IV. EXPERIMENTS

To validate the proposed framework for supervisory teleoperated control and local autonomy, we design a set of experiments addressing two complementary objectives.

First, we evaluated the effectiveness of the supervisory teleoperation scheme. In this setting, we investigate whether the multi-agent system can successfully and efficiently traverse a sequence of via-regions provided by the human operator in an obstacle-free environment, while maintaining minimal tracking error with respect to the specified trajectory setpoints.

Second, we assessed the performance of the proposed local autonomy and distributed control strategy in various cluttered environments. In this scenario, we examined whether the multi-agent system can reliably reach their respective goal configurations in the presence of dense obstacles, and we compared its performance with various gain configurations in terms of efficiency, robustness, and trajectory quality.




Fig. 3: A nominal affine formation $\mathcal{F}^* = (\mathcal{G}, \mathbf{P}^*)$ of the multi-robot system consisting of seven agents, including three leaders and four followers. The gray solid lines denote the undirected communication edges between neighboring agents defined by the graph topology $\mathcal{G}$. The coordinates associated with each agent represent their nominal positions in the formation configuration $\mathbf{P}^*$, where leader agent L_1 is defined as the origin of the nominal formation frame.

Across all experimental scenarios, we evaluated the multi-agent system consistently converged to the predefined nominal formation $\mathcal{F}^*$, as illustrated in Fig. 3, which is required for system calibration and coordination-based operation. All synthetic experiments are conducted using Python 3.10 on a workstation equipped with 13th Gen Intel Core™ i5-1335 CPU (12 threads) and 16.0 GiB of RAM, running on Ubuntu 22.04 (64-bit). Extensive and supplementary results are detailed in Appendix D.

A. Supervisory teleoperated control experiment

1) *Experimental Setup*: The supervisory teleoperation capability of the proposed framework was evaluated through a series of synthetic experiments conducted in an obstacle-free environment. The workspace dimensions were set to $100\text{ m} \times 100\text{ m}$, containing a robotic network composed of seven agents, of which three were designated as leaders and the remaining four as followers. The agents communicated through the predefined topology illustrated in Fig. 3. Each agent was assigned a communication radius of $r^c = 30\text{ m}$ with $\alpha^{\text{conn}} = 0.5$. Each agent has a physical radius of $r^r = 0.25\text{ m}$.

To emphasize the distinct configuration of leaders and followers within the framework, the control gains were selected accordingly. Leader agents were primarily tasked with tracking the prescribed reference trajectories, whereas follower agents focused on preserving the affine geometric formation. Consequently, the leaders were assigned trajectory-tracking weights $\mathbf{A}_{z,\text{leader}} = K_{\text{leader}}^{\text{DMP}} \mathbf{I}_D = 1050 \cdot \mathbf{I}_D$, while their formation-maintenance gains were disabled such that $\mathbf{A}_{\text{leader}}^{\text{form}} = \mathbf{B}_{\text{leader}}^{\text{form}} = \mathbf{0}_{2 \times 2}$. In contrast, the followers were assigned $\mathbf{A}_{z,\text{follower}} = \mathbf{0}_{2 \times 2}$ and formation-maintenance gains $\mathbf{A}_{\text{follower}}^{\text{form}} = \mathbf{B}_{\text{follower}}^{\text{form}} = D_{\text{follower}}^{\text{form}} \mathbf{I}_D = 1050 \cdot \mathbf{I}_D$. Both leaders and followers employed identical connectivity-preservation gains given by $\mathbf{A}^{\text{conn}} = K^{\text{conn}} \mathbf{I}_D = 1050 \cdot \mathbf{I}_D$,

and identical inter-agent collision avoidance gains of $\lambda = 1800, \beta = 2, \eta = 1$.

Each agent operated with a discrete control timestep of $\Delta t = 0.01$ s with base temporal scaling of $\tau = 10$, including an additional inter-agent temporal scaling parameters of $\tau_{\text{extra}}^{\text{agent}} = 50$ with distance threshold $d_{\text{slow}}^{\text{agent}} = 1$ m, $\tau_{\text{extra}}^{\text{form}} = 50$ with $d_{\text{slow}}^{\text{form}} = 15$ m, and $\tau_{\text{extra}}^{\text{conn}} = 50$ with $d_{\text{slow}}^{\text{conn}} = 15$ m. The navigation task was considered successfully completed once all the leader agents and follower agents reached their respective target positions within a positional and formation tolerance both of 10 cm, respectively. The target positions of the robotic network were determined adaptively based on the final target position and terminal heading of leader agent L_1 , together with the nominal affine formation configuration $\mathcal{F}^*$.

2) *Tasks*: The goal of the multi-agent robotic network is to navigate cooperatively through a sequence of prescribed via-regions before ultimately converging to a designated target destination. During the navigation process, the leader agents are responsible for tracking the reference trajectories and guiding the global motion of the formation, while the follower agents maintain the affine geometric configuration induced by the leaders. Throughout the task execution, all agents are additionally required to preserve safe inter-agent distances to prevent collisions and maintain sufficient communication connectivity.

To evaluate the supervisory teleoperation and trajectory-tracking capabilities of the proposed framework, two navigation scenarios with distinct trajectory geometries were considered. In both scenarios, the robotic network was initialized in the nominal affine formation configuration and tasked with traversing multiple via-regions before reaching the final destination.

In the first scenario, leader agent L_1 was assigned an *upside-down U-shaped* reference trajectory originating at (20, 20) m, passing through intermediate via-points located at (20, 80) m and (80, 80) m, and terminating at the target position (80, 10) m. The remaining leaders followed trajectories derived from the prescribed affine transformation, while the followers continuously adapted to preserve the desired formation geometry throughout the navigation task.

The second scenario was designed to further evaluate the robustness of the framework under trajectories with one additional via-region, larger curvature, and more abrupt directional changes. In this experiment, leader agent L_1 tracked an *S-shaped* reference trajectory originating at (30, 10) m, traversing via-points at (70, 10) m and (90, 50) m, (10, 50) m, and terminating at (50, 90) m. Similar to the first scenario, the remaining leaders followed the affine-transformed trajectories while the followers maintained the desired formation relative to the leaders.

For both navigation scenarios, the task was considered successfully completed once all agents reached their respective target positions within the prescribed positional tolerance while converging to the nominal formation configuration at the final destination. Throughout the execution, the followers were required to continuously adapt to geometric variations induced by the leaders, particularly during curved trajectory segments where transient tracking degradation caused temporary con-

traction of the formation geometry. Despite these variations, all agents were required to maintain collision-free navigation and preserve network connectivity for the duration of the task.

3) *Metrics*: The performance of the supervisory teleoperation framework is evaluated using the following metrics:

a) *Reference trajectory tracking error*: This metric evaluates how accurately the leader agents follow the reference trajectories defined by the sequence of via-regions. Let $z_i^{\text{ref},(k)}$ denote the reference position of agent i at time step k . The tracking error is defined as

$$e_{\text{traj}}^{(k)} = \sqrt{\frac{1}{N_L} \sum_{i \in \mathcal{V}_L} \|z_i^{(k)} - z_i^{\text{ref},(k)}\|_2^2}, \quad (24)$$

and averaged over time as

$$\bar{e}_{\text{traj}} = \frac{1}{T} \sum_{k=1}^T e_{\text{traj}}^{(k)}. \quad (25)$$

b) *Via-region tracking error*: This metric measures how well the formation passes through the specified via-regions. Let c_r denote the center of via-region r , and $\mathcal{T}_r$ the set of time steps during which the formation is expected to pass through it. The via-region error is defined as

$$e_{\text{via}}^{(r)} = \frac{1}{|\mathcal{T}_r|} \sum_{k \in \mathcal{T}_r} \sqrt{\frac{1}{N} \sum_{i=1}^N \|z_i^{(k)} - c_r\|_2^2}. \quad (26)$$

c) *Formation error*: The formation error quantifies the deviation of the followers' configuration from the desired formation:

$$e_{\text{form}}^{(k)} = \sqrt{\frac{1}{N_F} \sum_{i \in \mathcal{V}_F} \|z_i^{(k)} - z_i^{*,(k)}\|_2^2}, \quad (27)$$

with time-averaged value

$$\bar{e}_{\text{form}} = \frac{1}{T} \sum_{k=1}^T e_{\text{form}}^{(k)}. \quad (28)$$

d) *Safety margin*: The safety margin evaluates the minimum inter-agent distance throughout the execution phase:

$$M_{\text{safe}} = \min_k \min_{i \neq j} \left(\|z_i^{(k)} - z_j^{(k)}\|_2 - 2r^r \right), \quad (29)$$

where r^r denotes the physical radius of each agent. This means that a net-positive safety margin show no agents have collided throughout navigation. Moreover, the larger the value, the safer the agents navigate.

e) *Communication margin*: The communication margin quantifies the robustness of the underlying communication graph throughout the execution phase. For each pair of neighboring agents $(i, j) \in \mathcal{E}$, let $d_{ij}^{(k)} = \|z_i^{(k)} - z_j^{(k)}\|_2$ denote their distance at timestep k . The communication margin is defined as

$$\mu_{\text{comm}} = \min_k \min_{(i,j) \in \mathcal{E}} \left(r^c - d_{ij}^{(k)} \right). \quad (30)$$

Similar to (29), a net-positive safety margin means that all agents have not gone over the communication radius of its neighbors.

f) Computational complexity: The computational complexity metric evaluates the efficiency of the proposed method in terms of runtime and scalability with respect to the number of agents. To enable fair comparison across different system sizes, all runtime measures are normalized on a per-agent basis.

The average computation time per agent per simulation step is defined as

$$T_{\text{step}}^{\text{agent}} = \frac{1}{NT} \sum_{k=1}^T t^{(k)}, \quad (31)$$


where $t^{(k)}$ denotes the total wall-clock time required to compute control inputs for all agents at time step k , N is the number of agents, and T is the total number of time steps.

Similarly, the total runtime per agent for a complete simulation run is given by


$$T_{\text{run}}^{\text{agent}} = \frac{1}{N} \sum_{k=1}^T t^{(k)}. \quad (32)$$

4) Results: The spatial and temporal progression of the first scenario is illustrated in Fig. 4a, where subfigures (a)–(f) depict the sequential deployment of the robotic network alongside its evolving communication topology, subfigures (g) and (h) map the corresponding state trajectories along the x - and y -axes, and subfigures (i)–(l) detail the underlying performance metrics. Quantitatively, the robotic network successfully concluded the navigation profile within a total wall time of 3.43 s, executing 963 control cycles—which corresponds to an equivalent simulation time of 9.63 s. This temporal efficiency is underscored by a highly deterministic average computational overhead of 0.30 ms per agent per control loop. As shown in Fig. 4a(b), the network required approximately 1.93 s of simulation time to accelerate from rest and initiate meaningful spatial displacement. While this prolonged rise time is primarily a consequence of the large base time-scaling parameter utilized within the trajectory planner, such execution characteristics remain highly acceptable given the total trajectory length of approximately 198.7 m and the conservative nominal dynamics of the microrover platform, which operates at a baseline speed of 5.468 cm s^{-1} (as detailed in Appendix A).

Throughout this entire operation, system safety and network integrity were strictly preserved. All agents enforced robust collision avoidance and topological survival by maintaining a minimum safety margin of 2.165 m and a minimum connectivity margin of 18.58 m, respectively. Within these safe operational boundaries, the spatial accuracy of the system was driven by the tracking performance of the virtual leaders, which registered a mean reference trajectory tracking error of 20.08 cm. The peak tracking discrepancy localized near the via-point coordinate (80, 80) m at a maximum of 61.17 cm, demonstrating that the leader agents successfully navigated the prescribed via-regions while bounding trajectory drift well within sub-meter precision. Concurrently, the follower agents rigidly adhered to the leaders' tracking framework to preserve the global affine formation geometry, maintaining a




(a) Formation tracking state trajectories and chronological snapshots (a)–(f) outlining communication links, alongside time-history profiles along the x (g) and y (h) dimensions.




(b) Robotic network performance metrics: (i) trajectory and formation tracking errors for leader agents (L_1 – L_3) and follower agents (F_4 – F_7), (j) via-region clearance errors relative to the via-region radius (green dashed line), (k) worst-case safety margins relative to the violation threshold (red dashed line), and (l) computational execution workloads per control loop.

Fig. 4: Comprehensive experimental results and performance evaluation of the robotic network during the first navigation task. The mission profile for leader agent L_1 originates at (20, 20) m, navigates through via-points at (80, 20) m and (80, 80) m, and terminates at (80, 10) m. Detailed formation, safety, and computational metrics distributions are tracked via horizontal bars (median), triangles (mean), and circles (outliers).

tight mean formation tracking error of 3.480 cm. Although transient formation elongation and skewing were observed during periods of intense leader acceleration, owing to the inherent propagation delay required for the followers to synchronize with the positive velocity step, the nominal formation geometry was rapidly re-established once global velocity consensus was achieved. The coupled dynamics of AFC and the obstacle avoidance layer within this scenario are detailed



(a) Formation tracking state trajectories and chronological snapshots (a)-(f) outlining communication links, alongside time-history profiles along the x (g) and y (h) dimensions.



(b) Robotic network performance metrics: (i) trajectory and formation tracking errors for leader agents (L_1 - L_3) and follower agents (F_4 - F_7), (j) minimum closest approach errors to via-regions verifying boundary penetration (green dashed line), (k) worst-case safety margins relative to the violation threshold (red dashed line), and (l) execution workloads per control loop.

Fig. 5: Comprehensive experimental results and performance evaluation of the robotic network during the second navigation task. The mission profile for leader agent L_1 initiates at (30, 10) m, sweeps through via-points at (70, 10) m and (90, 50) m, (10, 50) m, and concludes at the target position (10, 50) m. Detailed formation, safety, and computation metrics distributions are explicitly tracked via horizontal bars (median), triangles (mean), and circles (outliers).

comprehensively in Appendix D.1.

The resulting navigation solution for the second scenario are illustrated in Fig. 5a, where subfigures (a)-(f) capture the temporal evolution of the robotic network, subfigures (g) and (h) map the corresponding state trajectories along the x - and y -axes, and subfigures (i)-(l) detail the underlying performance metrics. Quantitatively, the multi-agent network concluded the navigation profile in a wall time of 2.97 s,

executing 960 control loops, which corresponds to an equivalent simulated time of 96.4 s. This profile was supported by an average computational overhead of $232.8 \mu\text{s} \pm 0.45 \text{ ms}$ per control cycle. Global system stability was preserved within this extended space-time profile: all agents enforced robust collision avoidance and topological survival, maintaining a minimum communication margin of at least 8.458 m and a minimum inter-agent safety margin of 1.821 m. While the communication margin remained invariant relative to the first scenario, the safety baseline experienced a localized reduction driven by aggressive transient acceleration profiles out of the second via-region and sudden deceleration maneuvers approaching the third via-region, though it remained well within nominal operating tolerances.

Bounded by these safety guarantees, the internal system accuracy was dictated by the leader agents, which achieved a mean trajectory-tracking error of 30.22 cm. The peak tracking discrepancy was bounded at 1.33 m and localized near the via-point coordinate (10, 70) m. As observed in prior evaluations, the largest spatial deviations were strongly correlated with regions of high curvature. This behavior stems from the underlying C^1 continuity of the reference trajectories, which introduces discrete acceleration step changes, as visually evident near via-region (10, 70) in Fig. 5(e). These discontinuities momentarily saturated the ability of the DMP-based controller to ensure smooth, zero-lag tracking, inducing transient aggregation within the leader subset and a localized compression of the collective formation geometry. Nevertheless, the leaders consistently attenuated these disturbances, rapidly recovering accurate trajectory tracking following the transient deviations. This rapid recovery manifested as a sharp acceleration phase immediately after clearing the second via-region, as the DMP controllers aggressively nullified the tracking lag accumulated during the deceleration maneuvers through the first two curvatures. Concurrently, the follower agents tracked these leaders to preserve the desired global affine formation, registering an average formation tracking error of 5.61 cm. Although temporary geometric deformations occurred under extreme transient conditions, specifically where the leaders accelerated rapidly through highly curved segments and generated abrupt spatial variations, the followers successfully adapted within their kinematic constraints, rapidly attenuating the tracking lags to restore nominal formation integrity.

Overall, the empirical insights extracted from both experimental scenarios comprehensively validate the efficacy of the proposed robotic network architecture in executing complex, operator-specified cooperative navigation tasks. The leader agents consistently enforced path-tracking accuracy with sub-meter precision while rigorously satisfying all prescribed via-region constraints. Simultaneously, the follower agents regulated the global affine formation with high fidelity, demonstrating resilient adaptability despite transient geometric contractions and abrupt configuration changes induced by reference trajectory discontinuities. Most importantly, by preserving robust, bounded safety and communication margins across all test profiles, the architecture guarantees a reliable, collision-free, and topologically connected framework for multi-agent cooperative navigation in obstacle-free environment.

TABLE II: Number of obstacles per scenario [30]

Scenario	Rocks	Craters	Total obstacles
A	42	38	80
B	88	32	120
C	137	24	161

B. Local autonomy in cluttered environments experiment

1) *Experimental Setup*: The local autonomy capabilities of the proposed framework were evaluated through a set of synthetic simulations in highly cluttered environments. The experiments were conducted using a multi-agent system consisting of seven agents, comprising three leader agents and four follower agents. Throughout this experiment, all agents maintained a nominal configuration and interacted via the predefined communication topology illustrated in Fig. 3, establishing a baseline identical to the operation implemented in Experiment A.

The simulation environment was defined over a $120\text{ m} \times 120\text{ m}$ workspace and populated with static obstacles distributed within the convex region defined by the vertices $(10, 10)$, $(90, 10)$, $(5, 90)$, and $(90, 90)$. These obstacles consisted of rocks and craters of varying sizes. Their spatial distribution followed an exponential density model [30], parameterized by the maximum obstacle counts specified in Table II and constrained to a coverage density of approximately 15% of the convex hull area. This setup yields progressively more challenging navigation conditions across scenarios.

Each agent was modeled with a physical radius of $r^r = 0.25\text{ m}$ and equipped with a sensing radius of $r^s = 10\text{ m}$ to enable local perception of environmental obstacles. Inter-agent interaction was governed by a communication radius of $r^c = 30\text{ m}$, with the connectivity decay rate set to $\alpha^{\text{conn}} = 0.5$ and the connectivity gain parameterized as $\mathbf{A}^{\text{conn}} = K^{\text{conn}} \mathbf{I}_D$. To govern the system's temporal evolution, a base time-scaling factor of $\tau^{\text{base}} = 10$ was employed in conjunction with dedicated inter-agent, obstacle avoidance, formation tracking, and connectivity maintenance time-scaling parameters, all uniformly assigned as $\tau_{\text{extra}}^{\text{agent}} = \tau_{\text{extra}}^{\text{obs}} = \tau_{\text{extra}}^{\text{form}} = \tau_{\text{extra}}^{\text{conn}} = 50$.

To rigorously evaluate the performance and robustness of the proposed framework within highly cluttered environments, a Monte Carlo statistical procedure was implemented comprising 500 independent simulation runs per scenario under randomized obstacle layouts. Within this stochastic environment, the multi-agent network was systematically benchmarked against a baseline *pure DMP (decentralized)* framework to evaluate three alternative follower control configurations: *pure DMP (connected)*, *pure AFC*, and the proposed *hybrid DMP-AFC* formulation. The configuration-dependent follower gain matrices were defined using $\mathbf{A}_{z,\text{follower}} = K_{\text{follower}}^{\text{DMP}} \mathbf{I}_D$, $\mathbf{A}_{\text{follower}}^{\text{form}} = K_{\text{follower}}^{\text{form}} \mathbf{I}_D$, and $\mathbf{B}_{\text{follower}}^{\text{form}} = D_{\text{follower}}^{\text{form}} \mathbf{I}_D$, as detailed in Table III. Benchmarking directly against the decentralized DMP baseline isolated the trade-offs between destination reachability and formation maintenance under different control objectives, while the leader agents maintained a constant DMP gain of $\mathbf{A}_{z,\text{leader}} = K_{\text{leader}}^{\text{DMP}} \mathbf{I}_D = 250 \cdot \mathbf{I}_D$ across all trials to ensure a relaxed tracking baseline along the shortest path toward the destination.

TABLE III: Control gain configurations for the four follower control strategies.

Gain	Pure DMP (decent.)	Pure DMP (conn.)	Pure AFC	Hybrid DMP-AFC
$K_{\text{follower}}^{\text{DMP}}$	250	250	0	50
$K_{\text{follower}}^{\text{form}}$	0	0	300	250
$D_{\text{follower}}^{\text{form}}$	0	0	300	250
K^{conn}	0	1500	1500	1500

2) *Tasks*: The navigation task was defined such that the leader agents were required to navigate from their respective initial positions to designated target positions while avoiding collisions with both obstacles and neighboring agents. During task execution, the leaders implicitly governed the global motion of the robotic network, while the follower agents adapted their motions according to the evolving affine formation geometry, if AFC is applied.

At the beginning of each simulation run, leader agent L_1 was initialized at $(0, 0)\text{ m}$ and assigned a target position at $(100, 100)\text{ m}$. The initial and target positions of the remaining agents were determined according to the nominal formation configuration $\mathbf{P}^*$ illustrated in Fig. 3, with the group heading initialized at $\psi = \pi/4\text{ rad}$. Throughout the navigation process, the leader agents advance towards their target destinations and the follower agents proceed according to their set control gains, all agents were required to preserve communication connectivity by remaining within the communication radius of their neighboring agents, as well as maintaining safe distances with obstacles and other agents.

The task was considered successfully completed once all agents reached their respective target positions within a positional tolerance of $\varepsilon = 10\text{ cm}$. To ensure bounded simulation duration, each run was limited to a maximum of 100000 control steps. Any simulation run in which the agents failed to reach the prescribed goal configurations within the allocated control horizon was classified as a failed run.

3) *Metrics*: The performance of the proposed method is evaluated using several complementary metrics that quantify task success, communication robustness, and formation maintenance.

a) *Reachability*: Reachability evaluates the overall success rate of the system across multiple Monte Carlo runs per scenario. Let N_{run} denote the total number of simulation runs per scenario and N_{success} the number of runs per scenario in which all agents successfully reach their respective goal positions within a prescribed position tolerance $\varepsilon_{\text{run}} > 0 \in \mathbb{R}_+$ without deadlock or collision. The reachability is defined as

$$\nu = \frac{N_{\text{success}}}{N_{\text{runs}}} \quad (33)$$

A run is considered successful if, for all agents $i \in \mathcal{V}$,

$$\left\| \mathbf{z}_i^{(K)} - \mathbf{z}_i^{\text{goal}} \right\|_2 \leq \varepsilon_{\text{run}}. \quad (34)$$

This metric captures the robustness and reliability of the proposed method under randomized and highly cluttered environments.

b) *Path length*: The path length quantifies the efficiency of the trajectories executed by the agents. For agent i , the path length is defined as

$$L_i = \sum_{k=1}^{K-1} \left\| \mathbf{z}_i^{(k+1)} - \mathbf{z}_i^{(k)} \right\|_2, \quad (35)$$

where $\mathbf{z}_i^{(k)}$ denotes the position of agent i at discrete timestep k , and K is the total number of executed timesteps.

To evaluate system-level efficiency, the total path length across all agents is computed as

$$L = \sum_{i=0}^N L_i. \quad (36)$$

Lower path length indicates more efficient navigation with fewer detours caused by obstacle avoidance or inter-agent interactions.

c) *Formation error*: The formation error quantifying the deviation of the followers' configuration from the nominal formation, as previously described in (28).

d) *Communication margin*: The communication margin quantifies the robustness of the underlying communication graph throughout the execution phase. The communication margin is defined in (30).

e) *Safety margin*: The safety margin not only evaluate the inter-agent interactions, but also interaction with the obstacles. Let $\mathcal{O}$ denote the set of obstacles. The safety margin in (29) is expanded with obstacles and is defined as

$$\mu_{\text{safe}} = \min_k \left\{ \min_{i, m \in \mathcal{O}} \left\| \mathbf{z}_i^{(k)} - \mathbf{z}_m \right\|_2 - r^m, \min_{i \neq j} \left(\left\| \mathbf{z}_i^{(k)} - \mathbf{z}_j^{(k)} \right\|_2 - 2r^r \right) \right\}, \quad (37)$$

where r^m denotes the effective radius of obstacle m . Positive values indicate collision-free operation.

f) *Computational complexity*: The computational efficiency is evaluated by taking the mean total computation time over 500 runs per agent, using (31) and (32):

$$T_{\text{scenario}}^{\text{agent}} = \frac{1}{N_{\text{runs}}} \sum_{n=1}^{N_{\text{runs}}} \sum_{k=0}^{K-1} T_{\text{run}}^{\text{agent}}. \quad (38)$$

From an algorithmic perspective, each agent performs computations based on local interactions with its neighbors $\mathcal{N}_i$ and sensed obstacles $\mathcal{O}_i^{(k)}$. Thus, the per-agent computational complexity per step scales as $\mathcal{O}(|\mathcal{N}_i| + |\mathcal{O}_i^{(k)}|)$. Under the assumption of bounded communication degree and sensing radius, this results in approximately constant per-agent complexity, yielding overall linear scalability with respect to the number of agents.

4) *Results*: The quantitative performance of each control configuration across the three evaluation scenarios is summarized in Table IV, with the complete statistical distributions of each metric detailed in Appendix D.3 for a comprehensive evaluation. An analysis of these empirical data reveals distinct




Fig. 6: Trajectory evolution and time response of a 7-agent multi-agent system successfully navigating through a cluttered environment in scenario A using hybrid DMP-AFC approach. Panels (a)–(f) show representative snapshots of the agents' spatial configurations over time, including inter-agent communication links (gray lines) and current agent positions (markers). Panels (g) and (h) depict the temporal evolution of the agents' x - and y -coordinates, respectively.

trade-offs between the decentralized planning capability of DMPs and the structural rigidity offered by AFC.

Specifically, the *Pure DMP (decentralized)* framework yields the lowest mean total path lengths (e.g., 1153.70 m in Scenario C) and highly competitive execution times ranging from 2.031 s to 2.297 s. However, this trajectory efficiency comes at a severe cost to structural integrity and safety. It exhibits the highest formation errors across all scenarios, peaking at 5.0401 m in Scenario B, and registers dangerously negative communication margins down to -9.8159 m, which indicate frequent topology disconnections. Introducing explicit communication links in *Pure DMP (connected)* successfully rectifies connectivity, yielding the optimal communication margins of the study (≥ 4.7618 m) and the lowest overall wall times (≈ 1.95 s to 2.16 s). Nonetheless, bounded formation alignment remains poor at greater than 3.3 m.

Conversely, the *Pure AFC* configuration prioritizes tight topological tracking over spatial aggression. It achieves exceptional formation precision, maintaining an unparalleled low mean formation error (≤ 0.8295 m across all scenarios). This structural rigidity, however, forces the multi-agent system into highly conservative trajectories, resulting in a massive path length overhead (~ 1489 m to 1543 m) and nearly doubled wall times (≈ 3.8 s to 4.1 s) compared to its DMP counterparts. The proposed *Hybrid DMP-AFC* configuration successfully bridges this gap by leveraging the advantages of both methodologies. By combining the proactive trajectory generation of DMPs with the reactive tracking corrections of AFC, it achieves the highest overall mission success, leading in reachability across all evaluations (99.2%, 98.6%, and 99.0% for Scenarios A, B, and C, respectively). It simultaneously strikes a robust balance in formation error (1.80 m to 1.97 m) and guarantees a reliable, consistently positive communication margin (≥ 3.2472 m) without suffering the extreme path and

TABLE IV: Evaluation Metrics Across Configurations, Where A, B, and C Denote the Respective Evaluation Scenarios

Configuration	Reach. [%]			Path [m]			Form. Error [m]			Time [s]			Comm. Marg. [m]			Safety Marg. [m]		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
Pure DMP (decent.)	95.4	94.0	93.6	1157.86	1163.60	1153.70	4.4440	5.0401	4.8702	2.031	2.297	2.260	-7.4158	-9.8159	-7.7549	0.5022	0.3215	0.3112
Pure DMP (conn.)	96.6	94.0	95.2	1157.14	1170.81	1167.95	3.3426	3.6290	3.6082	1.952	2.160	2.165	5.4049	4.7618	5.1441	0.4947	0.3803	0.3490
Pure AFC	98.2	97.2	96.0	1523.85	1543.22	1489.17	0.7205	0.7643	0.8295	3.821	4.071	4.182	1.1455	0.8202	1.6823	0.3344	0.2751	0.2552
Hybrid DMP-AFC	99.2	98.6	99.0	1351.79	1347.59	1315.35	1.8081	1.9257	1.9766	3.607	3.833	3.936	3.2472	3.7015	4.1240	0.3981	0.3260	0.3050




Fig. 7: Trajectory evolution and time response of a 7-agent multi-agent trapped in a local minimum (red) in scenario A using hybrid DMP-AFC approach. Panels (a)–(f) show representative snapshots of the agents' spatial configurations over time, including inter-agent communication links (gray lines) and current agent positions (markers). Panels (g) and (h) depict the temporal evolution of the agents' x - and y -coordinates, respectively.

time inflation seen in the Pure AFC baseline.

To visually inspect these multi-agent dynamics under the hybrid controller, a representative storyboard illustrating a successful and efficient solution sequence is depicted in Fig. 6. The behavioral trajectory directly mirrors the nominal statistical values reported in Table IV. Driven by the adaptive velocity term, the agents proactively avoid incoming obstacles well ahead of impending collisions while strictly maintaining a safe inter-agent clearance baseline. Concurrently, the AFC consensus protocol ensures that proper topology connectivity is enforced. Whenever a pair of agents approaches the communication boundary, the control law generates a smooth attractive force, nudging the vehicles to close the gap and keeping the group tightly compact. Finally, as the leaders converge on the designated target coordinates, their underlying DMP control inputs decay exponentially. This decay naturally diminishes the reference commands propagated to the followers, causing the entire formation to smoothly stabilize and settle within the strict target tolerance boundaries, taking approximately an additional 25s to achieve complete steady-state convergence. The corresponding adaptive temporal synchronization is demonstrated in Appendix D.2.

Conversely, a limitation of the framework to environmental

obstacles is illustrated through the failed navigation of L_3 (red) presented in Fig. 7. Here, mission failure is primarily driven by the emergence of local minima within the potential field landscape, which represents an inherent drawback of potential-field-based navigation methodologies. This susceptibility becomes critically pronounced when dense obstacle clusters aggregate into nonconvex geometric profiles. Such configurations can physically trap a routing agent, forcing the remaining unconstrained agents to halt or drop behind indefinitely to prevent a catastrophic breach of the global communication topology.

Nevertheless, the coupled nature of the control law also enables successful trap-escape behaviors. Despite localized entrapment within a non-convex obstacle basin, the active formation consensus and connectivity feedback terms can inject sufficient structural restoring forces to pull trapped follower agents out of the local minimum. However, this cooperative escape mechanism exhibits an asymmetric dependency on the agent's role. If a leader agent becomes entrapped, it cannot leverage formation-tracking internal forces and must rely solely on the net topology connectivity vector. Consequently, leader escape is guaranteed to fail if the concave boundary of the local minimum points in a direction approximately opposite to the net connectivity pull, cementing the operational boundaries of the autonomous framework.

Lastly, despite the proposed follower controllers maintaining net-positive mean communication and safety margins across all scenarios alongside proven stability and convergence properties, certain outliers still violate the connectivity and safety boundaries of the agents, as illustrated in Figs. 20–22. These boundary violations are primarily driven by a transient imbalance in the interaction between the forward-advancing control components (the DMP and AFC terms) and the reactive safety components (the obstacle avoidance, inter-agent collision avoidance, and connectivity terms). Consequently, addressing these localized edge cases through adaptive parameter tuning or dynamic gain scheduling remains an area requiring additional attention to further improve worst-case performance metrics.

V. DISCUSSION

This study demonstrates that DMPs and AFC can be effectively unified to enable formation-aware, adaptive, and cooperative navigation on the lunar surface. Motivated by the critical need to incorporate human-in-the-loop supervision without inflating operator workload, eliminate single points

of failure, and minimize reliance on explicit plant models, the proposed framework successfully balances high-level path guidance with low-level autonomous safety. Structurally, because the framework decouples trajectory generation from a platform’s specific differential constraints, it circumvents the need for explicit system modeling and can be deployed across heterogeneous robotic networks.

At the high-level of human-guided trajectory tracking, empirical results validate that the leaders successfully steer the network through operator-specified via-regions with sub-meter precision, smoothly accommodating the acceleration discontinuities introduced by the C^1 continuity of the reference paths. Under these turning profiles, the follower agents rigidly preserve the global affine formation. However, a high-level behavioral artifact arises when leaders accumulate tracking lag: the subsequent exponential acceleration driven by the DMP tracking terms can induce transient structural deformations and skewing. Fortunately, because the AFC consensus protocol actively tracks both velocity and acceleration states, the collective formation rapidly attenuates these errors to restore nominal geometry. Furthermore, tuning the global temporal scaling factor τ^{base} offers a direct administrative mechanism to mitigate this lag by dynamically dampening the reference trajectory velocity.

Descending to the low-level mechanics of autonomous navigation, safety is enforced through volumetric artificial potential fields that dynamically reject multi-scale obstacles and avoid inter-agent collisions. By integrating the agents’ instantaneous velocity vectors into the potential function, the framework generates proactive, early avoidance maneuvers that scale with the sensing radius r^s . The inclusion of adaptive temporal scaling significantly improves this navigation capability. By dynamically modulating the temporal factor, it enables agents with high velocity or acceleration to actively slow down before reaching obstacles or neighboring agents, thereby facilitating smoother, more controlled deceleration profiles. Despite this predictive capability, extreme spatial discontinuities or *positional jumps* can still manifest if an agent penetrates an obstacle’s nominal safety radius. While saturating the peak repulsive force suppresses these severe control inputs, force saturation risks causing a total safety violation. Under dense obstacle configurations, the saturated avoidance vector can be dominated by competing forward-advancing or formation-holding forces, creating a net force vector that inadvertently drives the agent into the obstacle.

This low-level interaction between environmental avoidance and structural constraints exposes a systemic network limitation: the connectivity maintenance framework can inadvertently reintroduce a single point of failure within an otherwise decentralized system. If an individual agent becomes entrapped within a non-convex obstacle local minimum, the remaining unconstrained agents will enter a deadlock state, halting forward progress to prevent a breach of the communication topology radius r^c . This vulnerability is critically severe when a leader agent becomes trapped, given that a minimum of three operational leaders is structurally required to satisfy the affine formation constraints. To mitigate this, future iterations must either integrate a dynamic topology-switching failsafe

mechanism where a free follower automatically inherits the role of an entrapped leader, or implement a DMP-switching mechanism that transitions the network to a new dynamic movement primitive specifically tailored to escape local minima.

Furthermore, statistical outliers reveal connectivity boundary violations across various gain configurations, particularly within the pure AFC baseline. Because the baseline connectivity function omits higher-order velocity and acceleration feedback, the formation maintenance terms dominate the control loop as the formation scaling factor expands. Consequently, pure AFC lacks an intrinsic mechanism to guarantee network bounding when leaders diverge aggressively around obstacles. In the connected DMP and hybrid configurations, these residual violations are exacerbated by high-gain obstacle potential fields that deflect peripheral agents away from the cluster. Resolving these edge cases requires a structural upgrade to the connectivity controller, specifically, incorporating matching velocity and acceleration terms to balance out the AFC expansion forces, alongside multi-objective optimization to balance safety and connectivity gains.

When mapping these continuous control laws onto physical hardware that relies on discrete actuation primitives, such as the Lunar Zebro rovers, additional algorithmic modifications are required to translate continuous command velocities into sets of discrete motion primitives. From an online computational perspective, however, the proposed architecture remains highly feasible, matching the order of complexity of standard baseline methods. The primary bottleneck within the control loop stems from evaluating the volumetric obstacle avoidance, which scales linearly with the number of obstacles populated within the local sensing radius. Nevertheless, because the execution overhead remains strictly within the microsecond range per cycle, the global mission planning time is kept to a few seconds. This performance is highly efficient for slow-moving lunar rovers that displace at centimeter-per-second scales, confirming that the algorithm is fully capable of real-time, online deployment on resource-constrained flight computers prioritized for robustness and reliability.

Ultimately, assuming an obstacle-free destination zone and appropriately bounded collision avoidance parameters, DMP guarantees asymptotic convergence of the leader agents toward their target configurations, which inherently drives the follower agents toward its stable steady-state formation. While AFC performs exceptionally well within obstacle-free environments, its efficacy in maintaining sub-meter formation precision degrades in highly cluttered terrain. This performance degradation is primarily caused by the decentralized nature of the architecture, which leaves the leaders with no intrinsic knowledge of the global formation state when executing independent obstacle avoidance maneuvers. To address this limitation, additional DMPs could be integrated to enforce structural constraints, such as compelling the sub-leaders ($L_2, \dots, L_{N_L}$) to maintain a fixed relative offset from the primary leader (L_1) to prevent formation collapse. Alternatively, providing the leaders with real-time feedback regarding the collective configuration would enable continuous, shape-preserving adaptations that strictly adhere to admissible affine transformations.

For future work, several critical directions will advance the capabilities and technical readiness of this architecture. First, physical field validation will transition these control laws from synthetic environments onto physical multi-rover hardware platforms, such as the Lunar Zebro, evaluated within high-fidelity lunar analogue testbeds. Second, semantic hazard integration enhances the high-level supervisory teleoperation layer by embedding operator-defined, terrain-aware risk regions directly into the proactive trajectory generator, thereby preventing the multi-agent network from encroaching upon treacherous topography such as steep crater rims. Third, a decentralized, UWB-based relative localization scheme will be developed by adapting the affine control laws to utilize local barycentric coordinates referenced to a primary anchor node (L_1), eliminating reliance on global positioning. Finally, to transcend the heuristic limitations of potential fields, CBFs synthesized from the system's internal dynamics will be implemented to provide mathematically rigorous, hard safety guarantees for both inter-agent collision avoidance and topological connectivity maintenance.

REFERENCES

- [1] Y. Gao and S. Chien, "Review on space robotics: Toward top-level science through space exploration," *Science Robotics*, vol. 2, no. 7, p. eaan5074, 2017.
- [2] A. Abdel-Rahman, C. Cameron, B. Jenett, M. Smith, and N. Gershenfeld, "Self-replicating hierarchical modular robotic swarms," *Communications Engineering*, vol. 1, no. 1, p. 35, Nov 2022. [Online]. Available: <https://doi.org/10.1038/s44172-022-00034-3>
- [3] B. Martinez Rocamora, C. Kilic, C. Tatsch, G. A. S. Pereira, and J. N. Gross, "Multi-robot cooperation for lunar in-situ resource utilization," *Frontiers in Robotics and AI*, vol. Volume 10 - 2023, 2023. [Online]. Available: <https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2023.1149080>
- [4] M. J. Schuster, M. G. Müller, S. G. Brunner, H. Lehner, P. Lehner, R. Sakagami, A. Dömel, L. Meyer, B. Vodermayr, R. Giubilato, M. Vayugundla, J. Reill, F. Steidle, I. von Barga, K. Bussmann, R. Belder, P. Lutz, W. Stürzl, M. Smíšek, M. Maier, S. Stoneman, A. F. Prince, B. Rebele, M. Durner, E. Staudinger, S. Zhang, R. Pöhlmann, E. Bischoff, C. Braun, S. Schröder, E. Dietz, S. Frohmann, A. Börner, H.-W. Hübers, B. Foing, R. Triebel, A. O. Albu-Schäffer, and A. Wedler, "The arches space-analogue demonstration mission: Towards heterogeneous teams of autonomous robots for collaborative scientific sampling in planetary exploration," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5315–5322, 2020.
- [5] W. De Meyere, A. Shanbhag, and A. Menicucci, "Floating gate dosimeter characterization for space applications," *IEEE Transactions on Nuclear Science*, vol. 71, no. 8, pp. 1854–1863, 2024.
- [6] R. U. Sonsalla, S. Planthaber, R. Dominguez, A. Dettmann, F. Cordes, B. Huelsen, C. Schulz, P. Schoeberl, S. Kasperski, H. Wiedemann, and F. Kirchner, "Towards a semi-autonomous robotic exploration of a lunar skylight cavity," in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 1–20.
- [7] A. Hidding, H. Bier, L. Peternel, A. J. Becoy, F. A. P. Romio, and G. Calabrese, "Robotic lava tube mapping and multimodal data collection using quadruped and lidar," *Available at SSRN 5238812*, 2025.
- [8] A. Gautam and S. Mohan, "A review of research in multi-robot systems," in *2012 IEEE 7th International Conference on Industrial and Information Systems (ICIIS)*, 2012, pp. 1–5.
- [9] A. Banerjee, M. Mukherjee, S. Satpute, and G. Nikolakopoulos, "Resiliency in space autonomy: a review," *Current Robotics Reports*, vol. 4, no. 1, pp. 1–12, Mar 2023. [Online]. Available: <https://doi.org/10.1007/s43154-023-00097-w>
- [10] S. T. Arzo, D. Sikeridis, M. Devetsikiotis, F. Granelli, R. Fierro, M. Esmaili, and Z. Akhavan, "Essential technologies and concepts for massive space exploration: Challenges and opportunities," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 1, pp. 3–29, 2023.
- [11] J. C. Mankins, "Technology readiness assessments: A retrospective," *Acta Astronautica*, vol. 65, no. 9, pp. 1216–1223, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576509002008>
- [12] G. S. Aglietti, "Current challenges and opportunities for space technologies," *Frontiers in Space Technologies*, vol. 1, 2020. [Online]. Available: <https://www.frontiersin.org/journals/space-technologies/articles/10.3389/frspt.2020.00001>
- [13] T. Fong, J. Rochlis Zumbado, N. Currie, A. Mishkin, and D. L. Akin, "Space telerobotics: unique challenges to human-robot collaboration in space," *Reviews of Human Factors and Ergonomics*, vol. 9, no. 1, pp. 6–56, 2013.
- [14] T. Sheridan, "Space teleoperation through time delay: review and prognosis," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 592–606, 1993.
- [15] P. Schmaus, D. Leidner, T. Krüger, R. Bayer, B. Pleintinger, A. Schiele, and N. Y. Lii, "Knowledge driven orbit-to-ground teleoperation of a robot coworker," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 143–150, 2020.
- [16] V. Verma, M. W. Maimone, D. M. Gaines, R. Francis, T. A. Estlin, S. R. Kuhn, G. R. Rabideau, S. A. Chien, M. M. McHenry, E. J. Graser, A. L. Rankin, and E. R. Thiel, "Autonomous robotics is driving perseverance rover's progress on mars," *Science Robotics*, vol. 8, no. 80, p. eadi3099, 2023. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.adi3099>
- [17] E. H. J. Riemens and R. Thilak Rajan, "Distributed Navigation with Dynamic Obstacles," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [18] P. Mestres, C. Nieto-Granda, and J. Cortés, "Distributed Safe Navigation of Multi-Agent Systems Using Control Barrier Function-Based Controllers," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6760–6767, 2024.
- [19] Y. Yan, X. Li, X. Qiu, J. Qiu, J. Wang, Y. Wang, and Y. Shen, "Relative Distributed Formation and Obstacle Avoidance with Multi-agent Reinforcement Learning," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1661–1667.
- [20] M. Feng and H. Xu, "Multi-robot enhanced manet routing with situation aware online reinforcement learning," in *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2018, pp. 1145–1150.
- [21] P.-Y. Lajoie and G. Beltrame, "Swarm-slam: Sparse decentralized collaborative simultaneous localization and mapping framework for multi-robot systems," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 475–482, 2024.
- [22] Z. Lin, L. Wang, Z. Chen, M. Fu, and Z. Han, "Necessary and sufficient graphical conditions for affine formation control," *IEEE Transactions on Automatic Control*, vol. 61, no. 10, pp. 2877–2891, 2016.
- [23] Z. Han, K. Guo, L. Xie, and Z. Lin, "Integrated relative localization and leader-follower formation control," *IEEE Transactions on Automatic Control*, vol. 64, no. 1, pp. 20–34, 2019.
- [24] R. T. Rajan and A.-J. van der Veen, "Joint motion estimation and clock synchronization for a wireless network of mobile nodes," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012, pp. 2845–2848.
- [25] M. A. Shalaby, C. C. Cossette, J. R. Forbes, and J. Le Ny, "Calibration and uncertainty characterization for ultra-wideband two-way-ranging measurements," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 4128–4134.
- [26] D. Zhou, Z. Wang, and M. Schwager, "Agile Coordination and Assistive Collision Avoidance for Quadrotor Swarms Using Virtual Structures," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 916–923, 2018.
- [27] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: Learning attractor models for motor behaviors," *Neural Computation*, vol. 25, no. 2, pp. 328–373, 02 2013. [Online]. Available: https://doi.org/10.1162/NECO_a_00393
- [28] J. Umlauf, D. Sieber, and S. Hirche, "Dynamic movement primitives for cooperative manipulation and synchronized motions," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 766–771.
- [29] J. M. Cloud, W. J. Beks, and J. M. Schuler, "Vision-based movement primitives for lunar hazard avoidance," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 5142–5149.
- [30] T. Manteaux, D. Rodríguez-Martínez, and R. T. Rajan, "Rapf: Efficient path planning for lunar microrovers," in *2024 International Conference on Space Robotics (iSpaRo)*, 2024, pp. 56–63.

- [31] Z. Li, G. Leus, and R. T. Rajan, "Fast multiagent formation stabilization with sparse universally rigid frameworks," in *2025 33rd European Signal Processing Conference (EUSIPCO)*, 2025, pp. 2392–2396.
- [32] S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert, "Learning movement primitives," in *Robotics Research. The Eleventh International Symposium*, P. Dario and R. Chatila, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 561–572.
- [33] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921889008001772>
- [34] C. Yuksel, S. Schaefer, and J. Keyser, "Parameterization and applications of catmull-rom curves," *Computer-Aided Design*, vol. 43, no. 7, pp. 747–755, 2011, the 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010448510001533>
- [35] D.-H. Park, H. Hoffmann, P. Pastor, and S. Schaal, "Movement reproduction and obstacle avoidance with dynamic movement primitives and potential fields," in *Humanoids 2008 - 8th IEEE-RAS International Conference on Humanoid Robots*, 2008, pp. 91–98.
- [36] H. Hoffmann, P. Pastor, D.-H. Park, and S. Schaal, "Biologically-inspired dynamical systems for movement generation: Automatic real-time goal adaptation and obstacle avoidance," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 2587–2592.
- [37] M. Saveriano, F. J. Abu-Dakka, A. Kramberger, and L. Peternel, "Dynamic movement primitives in robotics: A tutorial survey," *The International Journal of Robotics Research*, vol. 42, no. 13, pp. 1133–1184, 2023.
- [38] S. Zhao, "Affine formation maneuver control of multiagent systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4140–4155, 2018.
- [39] M. Ginesi, D. Meli, A. Roberti, N. Sansonetto, and P. Fiorini, "Dynamic movement primitives: Volumetric obstacle avoidance using dynamic potential functions," *Journal of Intelligent & Robotic Systems*, vol. 101, no. 4, p. 79, Mar 2021. [Online]. Available: <https://doi.org/10.1007/s10846-021-01344-y>
- [40] 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.
- [41] N. E. Manitara and C. N. Hadjicostis, "Distributed stopping for average consensus in digraphs," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 3, pp. 957–967, 2018.
- [42] T. Maliappis, "Current state of the iac rover's obc locomotion software," Jul. 2023, version 1.3, IAC Rover.
- [43] S. Bandyopadhyay, J. Lazio, A. Stoica, P. Goldsmith, B. Blair, M. Quadrelli, J.-P. de la Croix, and A. Rahmani, in *2018 IEEE Aerospace Conference*, 2018, pp. 1–10.
- [44] J.-S. Lee, Y.-W. Su, and C.-C. Shen, "A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi," in *IECON 2007 - 33rd Annual Conference of the IEEE Industrial Electronics Society*, 2007, pp. 46–51.
- [45] H. Obeidat, W. Shuaib, O. Obeidat, and R. Abd-Alhameed, "A review of indoor localization techniques and wireless technologies," *Wireless Personal Communications*, vol. 119, no. 1, pp. 289–327, 7 2021. [Online]. Available: <https://doi.org/10.1007/s11277-021-08209-5>
- [46] M. Drobzyk, C. Strowik, and C. Philpot, "A wireless communication and positioning experiment for the iss based on ir-uwb," in *2017 IEEE Wireless Communications and Networking Conference (WCNC)*, 2017, pp. 1–6.
- [47] J. Ni, D. Arndt, P. Ngo, C. Phan, and J. Gross, "Uwb tracking system design for lunar/mars exploration," NASA Johnson Space Center, Houston, Texas, USA, Tech. Rep. NASA/TM-2008-214000, 2008. [Online]. Available: <https://ntrs.nasa.gov/citations/20080030264>
- [48] N. Panagiotopoulos, L. N. Daniel, O. Ratiu, and E. Umit, "Applications of ultra wide band technology in lunar exploration: An uwb lunar mission demonstration," in *2025 European Data Handling Data Processing Conference (EDHPC)*, 2025, pp. 1–6.
- [49] Qorvo, "Aps014 application note - antenna delay calibration of dw1000-based products and systems revision 1.3," 2024, last accessed 16 November 2025. [Online]. Available: <https://www.qorvo.com/products/p/DWM1001-DEV#documents>
- [50] —, "Aps011 application note - sources of error in dw1000-based two-way ranging (twr) schemes revision 1.2," 2024, last accessed 16 November 2025. [Online]. Available: <https://www.qorvo.com/products/p/DWM1001-DEV#documents>
- [51] —, "Product datasheet: Dwm1001-dev, dwm1001 module development board, revision 1.4," 2017, last accessed 16



Fig. 8: A Lunar Zebro microrover, TRAICI variant

- November 2025. [Online]. Available: <https://www.qorvo.com/products/p/DWM1001-DEV#documents>
- [52] R. T. Rajan and A.-J. van der Veen, "Joint ranging and clock synchronization for a wireless network," in *2011 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2011, pp. 297–300.
- [53] H. Wymeersch, S. Marano, W. M. Gifford, and M. Z. Win, "A machine learning approach to ranging error mitigation for uwb localization," *IEEE Transactions on Communications*, vol. 60, no. 6, pp. 1719–1728, 2012.
- [54] F. Shan, H. Huo, J. Zeng, Z. Li, W. Wu, and J. Luo, "Ultra-wideband swarm ranging protocol for dynamic and dense networks," *IEEE/ACM Transactions on Networking*, vol. 30, no. 6, pp. 2834–2848, 2022.
- [55] J. Cano, G. Pagès, Chaumette, and J. LeNy, "Clock and power-induced bias correction for uwb time-of-flight measurements," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2431–2438, 2022.
- [56] M. A. Shalaby, C. C. Cossette, J. R. Forbes, and J. Le Ny, "Calibration and uncertainty characterization for ultra-wideband two-way-ranging measurements," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 4128–4134.
- [57] A. Haggengruber, M. Krogus, and E. Olson, "Non-parametric error modeling for ultra-wideband localization networks," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 2568–2574.
- [58] S. Hayati, R. Volpe, P. Backes, J. Balaram, and W. Welch, "Microrover research for exploration of mars," in *Proceedings of the 1996 AIAA Forum on Advanced Developments in Space Robotics*, University of Wisconsin, Madison, 1996.
- [59] K.-M. Cheung, W. W. Jun, S. Bhamidipati, and P. Carter, "Ground-assisted position navigation and timing (pnt) for moon and mars," in *2024 IEEE Aerospace Conference*, 2024, pp. 1–19.
- [60] C. Sarkar, R. V. Prasad, R. T. Rajan, and K. Langendoen, "Sleeping beauty: Efficient communication for node scheduling," in *2016 IEEE 13th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, 2016, pp. 56–64.
- [61] S. M. Kay, "Statistical signal processing: estimation theory," *Prentice Hall*, vol. 1, pp. Chapter–8.8, 1993.
- [62] J.-P. Williams, D. Paige, B. Greenhagen, and E. Sefton-Nash, "The global surface temperatures of the moon as measured by the diviner lunar radiometer experiment," *Icarus*, vol. 283, pp. 300–325, 2017, lunar Reconnaissance Orbiter - Part II. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0019103516304869>

APPENDIX A LUNAR ZEBRO MICROROVER PLATFORM

The six-wheeled legged (whegged) rover platform, as illustrated in Fig. 8 utilized for system deployment is composed

of an integrated computing, sensing, and actuation architecture designed for multi-agent operations.

A. Computational Architecture and Power Distribution

The onboard computational core is a Raspberry Pi 5 single-board computer featuring an ARM quad-core Cortex-A76 CPU, 8 GB of RAM, and a 128 GB solid-state microSD card for data logging and storage. The system runs a real-time optimized Ubuntu 24.04 LTS operating system. High-level algorithmic execution, multi-agent communication, and sensor orchestration are handled via the Robot Operating System 2 (ROS 2) Jazzy. Power distribution to the computational sub-units and the actuation systems is managed via a 14.4 V lithium-ion battery pack with a nominal capacity of 4900 mA h, providing a stable power supply for short experimental field runs.

B. Sensor Suite and Localization Modalities

The perception and localization network combines optical, inertial, and radio-frequency modalities to enable accurate decentralized state estimation:

- **Vision System:** Primary exteroceptive sensing is performed by a dual optical camera array. Each imaging module incorporates a Sony IMX219 image sensor with a focal length of 3.04 mm and a horizontal field-of-view (FOV) of 62.2°. The sensors provide a maximum photo resolution of 3230×2462 pixels for localized spatial tracking.
- **Inertial Measurement Unit (IMU):** Kinematic data and local magnetic variations are captured via an Adafruit 9-degree-of-freedom (DOF) BNO055 intelligent IMU. This sensor integrates three primary triaxial components to fulfill distinct navigation and estimation roles:
 - **Accelerometer:** Measures total translational acceleration, which encompasses both the dynamic linear motion vectors of the rover and the static gravitational vector, in ms^{-2} . By decoupling these components, it provides independent metrics for linear acceleration and gravity direction, allowing the system to estimate inclination (pitch and roll) relative to the Earth's gravitational field.
 - **Gyroscope:** Measures the angular velocity vector across three orthogonal axes in rad s^{-1} . It detects rotational changes in real time, enabling the tracking of instantaneous orientation adjustments and transient maneuvering characteristics.
 - **Magnetometer:** Measures the ambient magnetic field strength vector in μT . By registering the local geomagnetic field vector, it serves as a digital compass that yields an absolute heading (yaw) reference, preventing long-term rotational drift inherent to gyroscopic integration.

An internal hardware fusion processor combines these complementary data streams to continuously synthesize absolute orientation vectors represented as both Euler angles and quaternions.

- **Ultra-Wideband (UWB) Localization:** Absolute 3D positioning data is obtained through an onboard Decawave DWM1001 UWB antenna module. This module calculates coordinates using time-of-flight (ToF) trilateration relative to static anchor nodes, establishing bounded spatial localization when the vehicle operates within the convex hull formed by the anchor network.

C. Actuation and Kinematic Modeling

The locomotion system consists of six independent, oblong circular-arc wheg segments optimized for traversing unstructured environments. Each individual wheg subtends an angle of 70° with an arc length of approximately 10.94 cm, yielding an effective wheel radius of $r_{\text{eff}} = 8.951$ cm.

Each wheg is actuated by a dedicated Faulhaber 1226S006B brushless DC (BLDC) motor operating within a nominal input voltage range of 6 V to 24 V. The motors deliver a rated continuous torque of up to 2.13 mNm, a stall torque of 10 mNm, and a no-load speed of up to 29 700 RPM. For torque amplification, each motor is coupled to a planetary gearhead with a 4 : 1 transmission reduction ratio. Primary low-level state feedback is supplied by integrated digital Hall-effect sensors functioning as incremental quadrature encoders.

The actuators are regulated by low-level, position-controlled Proportional-Integral-Derivative (PID) motor drivers that command each wheg to a target angular position $\vartheta \in S^1$. By modulating the target duration T over which a position trajectory sequence is executed, the average joint angular speed $\dot{\vartheta} = \omega$ is adapted. Through empirical verification, the minimum duration required to safely execute a full 360° joint rotation without undesired behaviors was determined to be $T^{\text{max}} = 2000$ ms. Consequently, the maximum achievable linear progression velocity of the rover is structurally bounded by:

$$v^{\text{max}} = \omega^{\text{max}} r_{\text{eff}} = \frac{70^\circ}{360^\circ} \frac{2\pi r_{\text{eff}}}{T^{\text{max}}} = 5.468 \text{ cm s}^{-1}. \quad (39)$$

D. Locomotion Control and Action Space

The mobility controller abstracts the underlying low-level joint synchronization into a discrete action space comprising $M = 8$ high-level motion primitives:

$$\mathcal{A} = \{\text{forward, backward, turn-left, turn-right, arc-fw-left, arc-fw-right, arc-bw-left, arc-bw-right}\}. \quad (40)$$

To command vehicle motion, each high-level primitive is explicitly parameterized with a scheduled execution duration, which maps directly to the execution period of a full gait sequence or single joint rotation [42]. This unified abstraction bridges the discrete physical actuation capabilities of the whegged hardware with continuous trajectory generation frameworks.

APPENDIX B

ULTRA-WIDEBAND CALIBRATION

A. Introduction

Space exploration is one of humanity's most ambitious endeavors, with robotic systems playing a crucial role in early extraterrestrial missions by advancing scientific understanding and preparing infrastructure for future human exploration [1]. Multi-robot systems are particularly advantageous in harsh environments such as the lunar surface, as they provide distributed sensing, redundancy, and scalability. Applications include radiation mapping [5], in-situ resource exploration [3], [4], lava tube detection for potential habitats [6], [7], and the deployment of large-scale infrastructures such as lunar telescope arrays for high-resolution astronomical observations [43].

To achieve large-area coverage in a cost-effective manner, fleets of low-end microrovers are deployed in cooperative networks spanning several meters. These rovers typically have a base mass below 10 kg [30], imposing strict constraints on power consumption, sensing, and computational resources. As coordination in such multi-agent systems relies on communication and localization, GNSS-denied environments such as the Moon require relative localization techniques [23]. Integrating communication and localization through wireless technologies reduces hardware complexity while maintaining functionality in resource-constrained settings. In particular, time-of-flight (TOF) measurements between rovers can support affine formation control (AFC) methods for relative localization using geometric constraints among agents [22].

Ultra-Wideband (UWB) communication represents a compelling solution for joint data exchange and precise localization in multi-agent networks, offering superior energy efficiency and distinct ranging precision compared to Wi-Fi and Bluetooth [44], [45]. This high-resolution timing capability has driven its adoption in aerospace applications, including object positioning aboard the International Space Station [46] and tracking extraterrestrial rovers in environments devoid of Global Navigation Satellite System (GNSS) infrastructure [47]. Furthermore, recent qualification campaigns have demonstrated that UWB hardware can successfully endure the rigorous environmental demands of spaceflight, validating its readiness for lunar exploration [48].

Despite these advantages, achieving high-fidelity UWB localization requires mitigating deterministic and stochastic error sources, such as antenna delay biases, clock desynchronization, and severe multipath propagation [49], [50]. Because these challenges scale with network size, multi-robot systems are often motivated to periodically assume known formations to systematically calibrate ranging biases and synchronize distributed internal clocks. While the practical operating range of low-power UWB transceivers is typically constrained to approximately 30 meters under line-of-sight (LOS) conditions, this envelope is well-suited to support the localized, cooperative maneuvers of microrover networks on the lunar surface.

[48] utilizes the DWM1000 module as the core radio frequency (RF) block for its lunar demonstration payload. This widely adopted UWB transceiver supports high-bandwidth




Fig. 9: Signal propagation model between transceiver nodes i and j . Blue boxes denote measured transmission ($\hat{T}_{ij}$, $\hat{T}_{ji}$) and reception ($\hat{R}_{ij}$, $\hat{R}_{ji}$) timestamps. The true TOF is τ_{ij} (red arrows). The total forward and reverse delays (yellow and red paths) are modeled as $\tau_{i \rightarrow j}$ and $\tau_{j \rightarrow i}$, where antenna-induced delays ($\mu_{(\cdot), TX}$, $\mu_{(\cdot), RX}$) are scaled by node-specific clock drift coefficients α_i and α_j .

data rates up to 6.8 Mbps [51]. However, when referenced to the physical layer, the timestamps recorded at the data-link layer exhibit a pronounced propagation delay driven by antenna-induced delays effects. This antenna-induced delay introduces a systematic timing offset, as illustrated in Figure 9. To mitigate the separate challenges of extreme radiation, vacuum, and thermal cycles in the lunar environment, the architecture presented in [48] embeds the DWM1000 within a hardened space-grade node, coupling it with a radiation-tolerant GR716 main processor unit (MCU), redundant flash memories, and dedicated thermal monitoring converters.

In addition to antenna-induced delays, each device operates using an independent internal clock. Variations arising from manufacturing tolerances, computational load, and environmental conditions can introduce discrepancies among these clocks [52]. Consequently, clock offsets and clock skews may arise even between nominally identical devices, which can significantly degrade localization accuracy. If left uncalibrated, this clock desynchronization compounding with the aforementioned antenna-induced delay propagates directly into the TOF calculations, introducing a severe ranging bias that scales as a function of the accumulating time offset. Therefore, to ensure coherent system operation and to enable real-time distributed relative localization for inter-agent coordination, inter-device calibration and precise clock synchronization are essential.

1) *Related Works:* Inter-device calibration and clock synchronization in UWB systems have evolved from physical-layer waveform adaptations to protocol-level optimizations and statistical learning frameworks. A central challenge across these domains is mitigating systematic, hardware-induced errors, such as internal clock drift, antenna delays, and signal-power dependencies, which severely degrade localization precision even in LOS conditions [53], [55].

Early literature prioritized mitigation directly at the physical layer. [53] mapped raw UWB waveforms to ranging bias estimates using non-parametric regression methods, providing rich error quantification at the cost of high computational demands. To address network scale, research shifted toward scalable, concurrent ranging and broadcast protocols [54]. While protocols like [54] minimize channel utilization via

TABLE V: A Review of State-of-the-Art UWB Calibration and Localization Solutions Compared to the Proposed Study.

Author (year)	Core Methodology	Error Sources Addressed	Infrastructure Dependence	Space Mission Ready
Wymeersch et al. (2012) [53]	Non-parametric ML Regression	Waveform distortion & biases	Large offline training dataset	No
Shan et al. (2022) [54]	Broadcast Timestamp Exchange	Protocol latency & scaling	High-density active nodes	No
Cano et al. (2022) [55]	Power-dependent Correction	Clock skew & path loss	Hardware-specific calibration	No
Shalaby et al. (2023) [56]	Linear LS with Cauchy Loss	Signal-dependent antenna delay	External motion capture system	No
Haggenmiller et al. (2024) [57]	Non-parametric Optimization	Empirical NLOS distributions	Known target network anchors	No
This study	Formation-constrained Joint Estimator	Coupled clock & antenna biases	None (Self-contained via geometry)	Yes

simultaneous single-to-many timestamp exchanges, their performance degrades under severe packet collisions in dense agent networks.

Concurrently, algebraic and statistical modeling of hardware-specific parameters gained traction. [55] leveraged lightweight, power-dependent models suitable for resource-constrained architectures, though the framework remains limited to LOS conditions. For multi-robot systems, [56] proposed a scalable antenna-delay calibration via a robust linear least-squares formulation with Cauchy loss, while [57] introduced a non-parametric optimization framework to jointly estimate hardware biases and node positions in non-line-of-sight (NLOS) environments. Similarly, decoupled joint localization and synchronization schemes have been studied to minimize structural coupling [52].

Despite these advancements, high-accuracy alternatives consistently rely on external ground-truth systems (e.g., motion capture) or decouple synchronization and calibration, which neglects residual coupling errors. Furthermore, these methods do not account for the operational boundaries of severe planetary environments lacking satellite navigation infrastructure [47], [48].

Table V outlines how the framework proposed in this work addresses these research gaps. In contrast to state-of-the-art solutions, our approach provides a coupled, resource-efficient estimation of both clock synchronization and antenna-induced delay biases to precisely calibrate the system. This onboard calibration then serves as a critical prerequisite input for multi-agent coordination strategies, such as AFC. By utilizing the calibrated nodes to establish rigid formations, the system can subsequently maintain exact inter-node baseline distances *a priori* and bypass the need for external positioning infrastructure.

2) *Contributions*: The contributions of this paper are summarized as follows:

- We propose a joint clock synchronization and antenna delay calibration tailored to low-end lunar microrovers operating in GNSS-denied environments.
- We validate the proposed approach through system-level analysis and experimental evaluation on real hardware, demonstrating reliable synchronization and ranging performance under realistic operating conditions and hardware imperfections.

- We release an open-source experimental dataset to support reproducibility and encourage further research in synchronization and ranging for resource-constrained multi-robot systems.

The remainder of this paper is organized as follows. Section B.2 presents the system model and requirements, together with background on offline and online joint calibration and ranging, as well as antenna-induced delays. Section B.3 describes the proposed method for estimating antenna-induced delays. Section B.4 presents the experimental evaluation and discussion of the main results. Finally, conclusions are drawn in Section B.5.

Notations: The Kronecker product is indicated by $\otimes$ and the transpose operator by $(\cdot)^T$. $\mathbf{1}_N = [1 \ 1 \ \dots \ 1] \in \mathbb{R}^{N \times 1}$ is a vector of ones and $\mathbf{I}_N$ is a $N \times N$ identity matrix.

B. Preliminaries

1) *System Modeling and Requirements*: We consider a distributed robotic network operating within a two-dimensional workspace $\mathbb{R}^2$, comprising a set of N homogeneous microrovers whose communication topology is modeled as an undirected, time-invariant graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Here, $\mathcal{V} = \{1, 2, \dots, N\}$ represents the vertex set of rovers and the edge set $\mathcal{E} = \{(i, j) \in \mathcal{V} \times \mathcal{V} \mid \|\mathbf{x}_i - \mathbf{x}_j\| \leq r^c, i \neq j\}$ defines the permissible bi-directional communication links within a uniform maximum transceiver range $r^c \in \mathbb{R}_{++}$. Crucially, we assume strict, bi-directional LOS conditions exist for every active edge $(i, j) \in \mathcal{E}$, meaning the UWB propagation paths are completely unobstructed by lunar terrain features, thereby eliminating NLOS stochastic attenuation. Furthermore, to perform the joint synchronization and antenna delay calibration, all rovers are assumed to remain fixed in place during the calibration sequence ($\dot{\mathbf{x}}_i = \mathbf{0} \ \forall i \in \mathcal{V}$). This node stationarity ensures that the true physical inter-node Euclidean distances $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$ remain constant and known *a priori*, providing the essential deterministic geometric constraints required to decouple and estimate co-dependent tracking biases before initiating dynamic multi-agent coordination.

Because the target platforms are localized, energy-constrained microrovers with restricted onboard capabilities, the system architecture must enforce minimal space complexity and adhere to strict computational bounds [58]. All

algorithms must execute deterministically in real time on resource-constrained embedded hardware, precluding the use of large memory-intensive data structures or high-power central processing units. Furthermore, the calibration framework must scale gracefully with an increasing number of network participants, ensuring that both communication overhead and processing demands remain tightly bounded as the rover swarm expands. To support prolonged lunar surface exploration, a baseline positional accuracy of 1 m (corresponding to a TOF resolution of approximately 3.34 ns) is established as sufficient for operational safety [59]. This requirement aligns with the sparse, open nature of the lunar landscape, which presents fewer tight obstructions compared to cluttered terrestrial environments. Consequently, because the platforms navigate at low velocities, the strict tracking demands on localization update frequency and sub-meter precision are relaxed, allowing the resource-constrained nodes to sustain safe navigation using coarser, computationally efficient estimation routines.

Mathematically, the network-wide estimation problem is framed around an explicit partitioning of environmental invariants and measured observables against hidden internal state variables. The known deterministic inputs to the calibration system consist of the static inter-node physical distances alongside the discrete sets of observed transmission and reception timestamps exchanged over active communication edges. Conversely, the unknown state vector to be jointly identified encompasses the internal clock skews, the absolute time clock offsets, and the non-negligible hardware antenna-induced delays unique between each rover pair.

2) *Offline Clock Synchronization and Ranging*: [52] proposed an algorithm for joint clock synchronization and distance estimation among stationary devices. Consider the two-way measurement process between a pair of devices, denoted as node i and node j . As illustrated in Figure 10, the devices exchange messages such that node i records the transmission timestamp $T_{ij,k}$ and the reception timestamp $R_{ij,k}$, while node j similarly records $T_{ji,k}$ and $R_{ji,k}$ during the k -th message exchange round.

The resulting pairwise two-way measurements can therefore be modeled as follows:

$$\alpha_i T_{ij,k} = \alpha_j R_{ji,k} - \beta_i + \beta_j - \tau_{ij} - \alpha_i q_{1,k} + \alpha_j q_{2,k}, \quad (41a)$$

$$\alpha_i R_{ij,k} = \alpha_j T_{ji,k} - \beta_i + \beta_j + \tau_{ij} - \alpha_i q_{3,k} + \alpha_j q_{4,k}. \quad (41b)$$

Here, α_i and β_i denote the encoded clock skew and clock offset parameters of node i , respectively, while α_j and β_j represent the corresponding parameters for node j . The quantity $\tau_{ij} \triangleq d_{ij}/c$ denotes the propagation delay between nodes i and j where d_{ij} is the pairwise distance between node i and j , and c is the speed of light in vacuum. q_k represents a statistical noise component at the k -th message.

For K message exchanges per node, the timestamps can be vectorized accordingly. At node i , the timestamp vector is defined as $\mathbf{t}_{ij} = [T_{ij,1}, R_{ij,1}, \dots, T_{ij,K}, R_{ij,K}]^T \in \mathbb{R}^{2K \times 1}$, while at node j , the corresponding vector is given by $\mathbf{t}_{ji} = [R_{ji,1}, T_{ji,1}, \dots, R_{ji,K}, T_{ji,K}]^T \in \mathbb{R}^{2K \times 1}$. Conse-




Fig. 10: Symmetric round-robin two-way TOF protocol between nodes (i, j) across K message exchanges. Nodes i and j log $K/2$ local transmission and reception timestamps, namely $T_{ij,k}$, $R_{ij,k}$ and $T_{ji,k}$, $R_{ji,k}$, respectively. The curved lines represent the nonlinear temporal evolution induced by time-varying clock drift and hardware instabilities (inspired by Fig. 1 in [52]).

quently, (41a) and (41b) can be extended to accommodate K message exchanges as follows:

$$\begin{bmatrix} \mathbf{t}_{ij} & -\mathbf{t}_{ji} & \mathbf{1}_{2K} & -\mathbf{1}_{2K} & \mathbf{e} \end{bmatrix} \begin{bmatrix} \alpha_i \\ \alpha_j \\ \beta_i \\ \beta_j \\ \tau_{ij} \end{bmatrix} = \mathbf{q}_{ij} \quad (42)$$

where $\mathbf{e} = [-1, +1, -1, +1, \dots]^T \in \mathbb{R}^{2K \times 1}$, and $\mathbf{q}_{ij}$ denotes the aggregated noise vector affecting the node pair. By selecting node i as the reference node, i.e., setting $\alpha_i = 1$ and $\beta_i = 0$, the vector $\mathbf{t}_{ij}$ can be moved to the right-hand side of the equation:

$$\mathbf{A}_{ji} \boldsymbol{\theta}_j = \mathbf{t}_{ij} + \mathbf{q}_{ij}, \quad (43)$$

where $\mathbf{A}_{ji} = [\mathbf{t}_{ji} \ \mathbf{1}_{2K} \ -\mathbf{e}] \in \mathbb{R}^{2K \times 3}$ and $\boldsymbol{\theta}_j = [\alpha_j \ \beta_j \ \tau_{ij}]^T \in \mathbb{R}^{3 \times 1}$. Moreover, since the resulting system matrix $\mathbf{A}_{ji}$ is fully ranked, the estimation problem admits a unique solution for the clock parameters at node j as well as for the propagation delay between nodes i and j . The parameters, $\boldsymbol{\theta}_j$, is obtained by minimizing the least squares norm:

$$\hat{\boldsymbol{\theta}}_j = \arg \min_{\boldsymbol{\theta}_j} \|\mathbf{A}_{ji} \boldsymbol{\theta}_j - \mathbf{t}_{ij}\|_2^2 = (\mathbf{A}_{ji}^T \mathbf{A}_{ji})^{-1} \mathbf{A}_{ji}^T \mathbf{t}_{ij}. \quad (44)$$

The optimization problem can be extended to a network comprising N nodes, represented as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. To obtain a globally optimal estimate of all $N - 1$ unknown clock parameters and all pairwise propagation delays using a single clock reference, the graph must be connected, ensuring that every node is reachable from the reference. By selecting node 1 as the reference, i.e., imposing $[\alpha_1, \beta_1] = [1, 0]$, and denoting the timestamp vector at node 1 by $\mathbf{t}_1$, network-wide synchronization and ranging can be formulated as the following least-squares minimization problem:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \|\mathbf{A} \boldsymbol{\theta} + \mathbf{t}_1\|_2^2 = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{t}_1, \quad (45)$$

where

$$\mathbf{A} = [\mathbf{T} \quad \mathbf{E}_1 \quad \mathbf{E}_2] \in \mathbb{R}^{2KM \times L},$$

$$\boldsymbol{\theta} = [\boldsymbol{\alpha} \quad \boldsymbol{\beta} \quad \boldsymbol{\tau}]^T \in \mathbb{R}^{L \times 1}.$$

$\mathbf{T} \in \mathbb{R}^{2KM \times N}$ aggregates the timestamp vectors from all $N - 1$ unknown nodes, while $\mathbf{E}_1 \in \mathbb{R}^{2KM \times N}$ is a matrix of linear combinations of $\pm \mathbf{1}_{2K}$ that encodes the clock offsets. The matrix $\mathbf{E}_2 = -\mathbf{I}_M \otimes \mathbf{e} \in \mathbb{R}^{2KM \times M}$ characterizes the pairwise propagation delays, with M denoting the number of node pairs, and $L = 2N + M - 2$. It is important to note that network connectivity is a prerequisite for feasible synchronization. However, missing links may lead to suboptimal estimation of pairwise distances between those particular nodes.

3) *Online Clock Synchronization and Ranging*: In energy-constrained and operational environments such as the lunar surface, minimizing power consumption and ensuring execution efficiency is critical, particularly when calibration procedures run online on low-end microrovers with strictly bounded processing and memory resources. To accommodate these hardware limitations, the block least-squares problem can be reformulated into a recursive, iterative least-squares approach [60]. This recursive modification compresses the space complexity from $\mathcal{O}(K)$ to $\mathcal{O}(1)$ for pairwise tracking by eliminating the need to store historical timestamp profiles.

Consider an active pairwise communication link between node i and node j , the sequential parameter estimation update upon receiving the k -th two-way measurement vector $\mathbf{t}_{ij,k} \in \mathbb{R}^{2 \times 1}$ is expressed as follows:

$$\mathbf{K}_{ji,k} = \boldsymbol{\Sigma}_{ji,k-1} \mathbf{A}_{ji,k}^T \left(\mathbf{I}_2 + \mathbf{A}_{ji,k} \boldsymbol{\Sigma}_{ji,k-1} \mathbf{A}_{ji,k}^T \right)^{-1} \quad (46a)$$

$$\hat{\boldsymbol{\theta}}_{j,k} = \hat{\boldsymbol{\theta}}_{j,k-1} + \mathbf{K}_{ji,k} (\mathbf{t}_{ij,k} - \mathbf{A}_{ji,k} \hat{\boldsymbol{\theta}}_{j,k-1}), \quad (46b)$$

$$\boldsymbol{\Sigma}_{ji,k} = (\mathbf{I}_3 - \mathbf{K}_{ji,k} \mathbf{A}_{ji,k}) \boldsymbol{\Sigma}_{ji,k-1} (\mathbf{I}_3 - \mathbf{K}_{ji,k} \mathbf{A}_{ji,k})^T \quad (46c)$$

where $\hat{\boldsymbol{\theta}}_{j,k} \in \mathbb{R}^{3 \times 1}$ denotes the updated parameter estimate vector at index k , $\boldsymbol{\Sigma}_{ji,k} \in \mathbb{R}^{3 \times 3}$ represents its associated estimation error covariance matrix, and $\mathbf{A}_{ji,k} \in \mathbb{R}^{2 \times 3}$ is the local measurement matrix mapping the timestamps logged by node j . $\mathbf{K}_{ji,k} \in \mathbb{R}^{3 \times 2}$ represents the estimator gain matrix, which dictates the weight assigned to the measurement innovation residual $(\mathbf{t}_{ij,k} - \mathbf{A}_{ji,k} \hat{\boldsymbol{\theta}}_{j,k-1})$.

For centralized network-wide clock synchronization and ranging with N many nodes, the formulation can be extended to

$$\mathbf{K}_k = \boldsymbol{\Sigma}_{k-1} \mathbf{A}_k^T (\mathbf{I}_{2M} + \mathbf{A}_k \boldsymbol{\Sigma}_{k-1} \mathbf{A}_k^T)^{-1}, \quad (47a)$$

$$\hat{\boldsymbol{\theta}}_k = \hat{\boldsymbol{\theta}}_{k-1} + \mathbf{K}_k (\mathbf{t}_{1,k} - \mathbf{A}_k \hat{\boldsymbol{\theta}}_{k-1}), \quad (47b)$$

$$\boldsymbol{\Sigma}_k = (\mathbf{I}_L - \mathbf{K}_k \mathbf{A}_k) \boldsymbol{\Sigma}_{k-1} (\mathbf{I}_L - \mathbf{K}_k \mathbf{A}_k)^T, \quad (47c)$$

where $\mathbf{A}_k, \mathbf{A}_{k-1} \in \mathbb{R}^{2M \times L}$ are the network-wide measurement matrices, and $\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}_{k-1} \in \mathbb{R}^{L \times L}$ are the associated estimator error covariance matrices at time instants k and $k-1$, respectively. By exploiting this recursive structure, the storage complexity is reduced from $\mathcal{O}(KLM)$ in the offline approach to $\mathcal{O}(LM)$, resulting in significantly improved memory efficiency without degrading estimation performance.

4) *Device-specific Antenna-Induced Delays*: In practice, an inherent delay exists between the measured timestamps and the effective antenna reference point due to hardware and design limitations. This delay introduces a systematic error in the observations that cannot be captured by the statistical noise term $\mathbf{q}$. As illustrated in Figure 9, the observed transmission and reception timestamps obtained from the data link layer at an arbitrary node, denoted by $\tilde{T}$ and $\tilde{R}$, respectively, can be expressed as

$$\tilde{T} = T - \mu_T, \quad (48a)$$

$$\tilde{R} = R + \mu_R. \quad (48b)$$

where μ_T and μ_R denote the antenna delays associated with the TX and RX channels of the considered node, respectively. For the sake of clarity, we first revisit the case of pairwise communication (41). The observed TOF, denoted by $\tilde{\tau}$, from node i to node j , denoted with subscript $i \rightarrow j$, and $j \rightarrow i$ vice versa, are given by:

$$\tilde{\tau}_{i \rightarrow j} = \tau_{ij} + \alpha_i \mu_{i,T} + \alpha_j \mu_{j,R}, \quad (49a)$$

$$\tilde{\tau}_{j \rightarrow i} = \tau_{ij} + \alpha_i \mu_{i,R} + \alpha_j \mu_{j,T}. \quad (49b)$$

where τ_{ij} denotes the true propagation delay between node i and node j , and $\mu_{i,T}$ and $\mu_{i,R}$ represent the TX and RX antenna delays of node i , respectively, and T and R denote the actual timestamps the signal leaves and arrives, respectively. Similarly, $\mu_{j,T}$ and $\mu_{j,R}$ denote the corresponding TX and RX antenna delays of node j . It is important to emphasize that these antenna delays are scaled according to the clock skew of their corresponding node. Moreover, $\tilde{\tau}_{i \rightarrow j}$ and $\tilde{\tau}_{j \rightarrow i}$ are not necessarily identical, since the antenna delays associated with the TX and RX channels at each incidental node may differ.

To quantify the systematic error introduced by the antenna delays, the timestamps in (41a) and (41b) must be substituted in (48a)-(48b). Since the systematic error in the estimated propagation delay represents the aggregate effect of the antenna delays introduced by the respective devices, the individual antenna delays relative to each device are not of interest. Instead, the focus lies on their combined contribution to the overall observation. Consequently, the clock skew factors can be absorbed into the antenna delay terms. This reformulation further preserves the convexity of the resulting optimization problem. For notational clarity, we define $\alpha_i \mu_{i,T} \triangleq \delta_{i,T}$ as the effective TX antenna delay at node i , and similarly denote $\delta_{i,R}$ as the effective RX antenna delay at node i . This procedure leads to the following expression:

$$\alpha_i \tilde{T}_{ij,k} - \alpha_j \tilde{R}_{ji,k} + \beta_i - \beta_j + \tau_{ij} = -\delta_{i,T} - \delta_{j,R} - \alpha_i q_{1,k} + \alpha_j q_{2,k}, \quad (50a)$$

$$\alpha_i \tilde{R}_{ij,k} - \alpha_j \tilde{T}_{ji,k} + \beta_i - \beta_j - \tau_{ij} = \delta_{i,R} + \delta_{j,T} - \alpha_i q_{3,k} + \alpha_j q_{4,k}. \quad (50b)$$

By substituting (49a) and (49b) into (50a) and (50b), respectively, we get

$$\begin{aligned} \alpha_i \tilde{T}_{ij,k} - \alpha_j \tilde{R}_{ji,k} + (\beta_i - \beta_j) + \tilde{\tau}_{i \rightarrow j} \\ = -\alpha_i q_{1,k} + \alpha_j q_{2,k}, \end{aligned} \quad (51a)$$

$$\begin{aligned} \alpha_i \tilde{R}_{ij,k} - \alpha_j \tilde{T}_{ji,k} + (\beta_i - \beta_j) - \tilde{\tau}_{j \rightarrow i} \\ = -\alpha_i q_{3,k} + \alpha_j q_{4,k}. \end{aligned} \quad (51b)$$

which demonstrates that intrinsic antenna delays affect only ranging, while clock synchronization remains unaffected. Accordingly, clock parameters between devices can still be reliably estimated by solving an unconstrained least-squares optimization problem in which the propagation delay, τ , is excluded from the parameter space, θ . Moreover, using affine model for clocks, the true propagation delay, τ , and the antenna delays, δ , cannot be jointly estimated, as each parameter depends on prior knowledge of the other. This interdependence leads to a lack of identifiability within the proposed model without additional constraints.

C. Proposed Algorithms

1) *Antenna-Induced Delay Estimation:* To estimate the systematic errors arising from antenna-induced delays at each node, the true propagation delay τ_{ij} must be isolated from the observed round-trip time measurements between nodes i and j , denoted by $\tilde{\tau}_{i \rightarrow j}$ and $\tilde{\tau}_{j \rightarrow i}$. In this work, the true propagation delays are assumed to be known during dedicated calibration phases, which are enabled by enforcing a known geometric configuration of the robotic team, such as an affine formation. This controlled formation provides prior knowledge of inter-agent distances, allowing the geometric propagation component to be separated from hardware-induced timing offsets.

However, this separation is complicated by the presence of unknown clock parameters across the network, which must be estimated prior to extracting antenna-induced effects. Specifically, the clock state of node i is parameterized as $\theta_i \triangleq [\alpha_i, \beta_i]$, and analogously for node j as θ_j . These parameters are estimated by reformulating (50a) and (50b), where the propagation distance terms are moved to the right-hand side to express the measurements as functions of the unknown clock parameters. This leads to a least-squares formulation over a set of $2K$ measurements, given as follows:

$$\hat{\theta}_j = \arg \min_{\theta_j} \|\bar{\mathbf{A}}_{ji} \bar{\theta}_j - \tilde{\mathbf{t}}_{ij}\|_2^2 = (\bar{\mathbf{A}}_{ji}^T \bar{\mathbf{A}}_{ji})^{-1} \bar{\mathbf{A}}_{ji}^T \tilde{\mathbf{t}}_{ij}, \quad (52)$$

where

$$\begin{aligned} \bar{\mathbf{A}}_{ji} &= [\tilde{\mathbf{t}}_{ji} \quad \mathbf{1}_{2K}] \in \mathbb{R}^{2KM \times 2(N-1)}, \\ \tilde{\mathbf{t}}_{ij} &= [\tilde{T}_{ij,1} \quad \tilde{R}_{ij,1} \quad \cdots \quad \tilde{T}_{ij,K} \quad \tilde{R}_{ij,K}]^T \in \mathbb{R}^{2K \times 1}, \\ \tilde{\mathbf{t}}_{ji} &= [\tilde{R}_{ji,1} \quad \tilde{T}_{ji,1} \quad \cdots \quad \tilde{R}_{ji,K} \quad \tilde{T}_{ji,K}]^T \in \mathbb{R}^{2K \times 1}. \end{aligned}$$

We compute the residual vector over the $2K$ measurements, denoted by ϵ_{ij} , by subtracting from the observed time vector $\tilde{\mathbf{t}}_{ij}$ the true propagation delay term $\tau_{ij} \triangleq \tau_{ij} \mathbf{e}_K \in \mathbb{R}^{2K \times 1}$

between nodes i and j . This formulation captures the aggregated effect of the estimated TX and RX antenna delays. The resulting residual vector, ϵ_{ij} , quantifies the discrepancy between the measurements and the model, and its statistical characteristics provide insight into potential systematic biases or model mismatches.

$$\epsilon_{ij} = \mathbf{A}_{ji} \hat{\theta}_j - (\tilde{\mathbf{t}}_{ij} - \tau_{ij}) = \begin{bmatrix} -\Delta_{i \rightarrow j} \\ \Delta_{j \rightarrow i} \end{bmatrix} \otimes \mathbf{1}_K \in \mathbb{R}^{2K \times 1}, \quad (53)$$

where $\Delta_{i \rightarrow j} = \delta_{i,T} + \delta_{j,R}$ and $\Delta_{j \rightarrow i} = \delta_{j,T} + \delta_{i,R}$, and $\delta_{i,T}$ and $\delta_{i,R}$ denote the estimated TX and RX antenna delays at node i , and similarly $\delta_{j,T}$ and $\delta_{j,R}$ at node j . Extending this to a network of N nodes, the least square solution for synchronizing the clocks with respect to node 1 becomes

$$\hat{\theta} = \arg \min_{\theta} \|\bar{\mathbf{A}} \bar{\theta} - \tilde{\mathbf{t}}_1\|_2^2 = (\bar{\mathbf{A}}^T \bar{\mathbf{A}})^{-1} \bar{\mathbf{A}}^T \tilde{\mathbf{t}}_1 \in \mathbb{R}^{2(N-1) \times 1}, \quad (54)$$

and the network-wide residuals over $2K$ measurements and M edges

$$\epsilon = \bar{\mathbf{A}} \hat{\theta} - (\tilde{\mathbf{t}}_1 - \tau) = \begin{bmatrix} -\Delta_{1 \rightarrow 2} \\ \Delta_{2 \rightarrow 1} \\ -\Delta_{1 \rightarrow 3} \\ \vdots \\ \Delta_{N \rightarrow N-1} \end{bmatrix} \otimes \mathbf{1}_K \in \mathbb{R}^{2KM \times 1}, \quad (55)$$

where $\tau \triangleq [\tau_{12}, \tau_{13}, \dots, \tau_{(N-1)N}] \otimes \mathbf{e} \in \mathbb{R}^{2KM \times 1}$ denotes the vector of true pairwise distances.

To validate the obtained results, we perform the original clock synchronization and ranging procedures by compensating the antenna delays and subsequently assess whether the estimated pairwise distances are consistent with the corresponding ground-truth distances. It is important to emphasize that node 1 cannot be adopted as the reference node for antenna delay compensation, i.e., by imposing $\delta_{1,T} \triangleq \delta_{1,R} \triangleq 0$. The observed timestamps acquired at node 1, denoted by $\tilde{\mathbf{t}}_1$, are likewise affected by antenna delays and therefore require compensation. Neglecting this effect would reintroduce bias in the estimation of pairwise distances during the joint processes of clock synchronization and ranging.

The individual antenna delays are estimated separately through a least-squares formulation

$$\hat{\delta} = \arg \min_{\delta} \|\mathbf{B} \delta - \epsilon\|_2^2 = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \epsilon \in \mathbb{R}^{2N \times 1}, \quad (56)$$

where

$$\delta = [\delta_{1,T} \quad \delta_{1,R} \quad \cdots \quad \delta_{N,T} \quad \delta_{N,R}]^T \in \mathbb{R}^{2N \times 1}, \quad (57)$$

$$\mathbf{B} = \left(\begin{bmatrix} \mathbf{B}_{i_1 j_1} \\ \mathbf{B}_{i_2 j_2} \\ \vdots \\ \mathbf{B}_{i_M j_M} \end{bmatrix} \otimes \mathbf{1}_K \right) \in \mathbb{R}^{2KM \times 2N}, \quad (58)$$

where each local block associated with the communication pair (i_m, j_m) is defined as

$$\mathbf{B}_{i_m j_m} = \begin{bmatrix} -\mathbf{e}_{2i_m-1}^T & -\mathbf{e}_{2j_m}^T \\ \mathbf{e}_{2i_m}^T & \mathbf{e}_{2j_m-1}^T \end{bmatrix} \in \mathbb{R}^{2 \times 2N}, \quad (59)$$

with $\mathbf{e}_n \in \mathbb{R}^{2N}$ denoting the n -th canonical basis vector. The first row of $\mathbf{B}_{i_m j_m}$ corresponds to the aggregated transmission-delay term $-\delta_{i_m, T} - \delta_{j_m, R}$ in (50a), whereas the second row corresponds to the reception-delay term $\delta_{i_m, R} + \delta_{j_m, T}$ in (50b). Consequently, the matrix $\mathbf{B}$ compactly represents the contribution of all antenna-delay parameters across the M communicating node pairs and K timestamp exchanges.

2) *Proposed Calibration Method*: Having established that antenna delays can be accurately estimated and subsequently removed from the observed timestamps, thereby enabling effective calibration of antenna-delay induced devices, it is of practical interest to implement the proposed procedure directly on the hardware. Such an implementation would facilitate autonomous calibration and ensure scalability as the network size increases.

To this end, the algorithm is executed in a centralized manner at a designated leader agent (node 1), which aggregates all timestamp measurements and performs joint calibration, clock synchronization, and antenna delay estimation within a unified optimization framework. In this formulation, node 1 assumes the role of a reference node in the classical synchronization sense, thereby resolving the inherent time-scale and time-offset ambiguities through the normalization constraints, without imposing any restrictions on the antenna delay parameters. This centralized architecture guarantees global consistency of all estimates, as the clock and delay parameters are intrinsically coupled across all node pairs and therefore require network-wide information for reliable identification. It is further assumed that all nodes maintain continuous connectivity with their neighbors and that pairwise inter-node distances are known *a priori* through formation maintenance, thereby preserving the network geometry over time.

To enable antenna delay compensation, the delay vector $(\hat{\delta})$ must first be estimated. However, its computation depends on prior knowledge of the clock parameters of all nodes relative to the reference node (node 1), denoted by $(\hat{\theta})$. To systematically address this interdependence, the estimation problem is therefore formulated within a constrained least squares (CLS) framework [61]. Consider a network consisting of N nodes, where node 1 serves as the reference. The network-wide calibration model can be expressed as

$$\mathbf{H}\psi = [\bar{\mathbf{A}} \quad \mathbf{B}] \begin{bmatrix} \bar{\theta} \\ \delta \end{bmatrix} = -\tilde{\mathbf{t}}_1 + \tau, \quad (60)$$

where $\psi = \begin{bmatrix} \bar{\theta}^T & \delta^T \end{bmatrix}^T$ collects the unknown clock and antenna delay parameters.

Importantly, the clock parameters are not affected by antenna delays and can therefore be estimated independently using an unconstrained least squares procedure. Once $\hat{\theta}$ is obtained, the antenna delays can be estimated in a subsequent

step by treating the clock parameter estimates as fixed quantities within the optimization problem. This sequential strategy preserves the structural separation between clock and delay parameters while ensuring consistent estimation.

where the constraint enforces consistency with the clock synchronization model. The resulting CLS solution decomposes into the following two subproblems:

$$\min_{\psi} \|(\tau - \tilde{\mathbf{t}}_1) - \mathbf{H}\psi\|_2^2 \quad \text{s.t.} \quad \bar{\mathbf{A}}\bar{\theta} = -\tilde{\mathbf{t}}_1, \quad (61)$$

where the constraint enforces consistency with the previously estimated clock parameters. The CLS solution becomes

$$\hat{\bar{\theta}} = \arg \min_{\bar{\theta}} \|\bar{\mathbf{A}}\bar{\theta} - \tilde{\mathbf{t}}_1\|_2^2 = (\bar{\mathbf{A}}^T \bar{\mathbf{A}})^{-1} \bar{\mathbf{A}}^T \tilde{\mathbf{t}}_1 \in \mathbb{R}^{2(N-1) \times 1} \quad (62)$$

$$\hat{\delta} = \arg \min_{\delta} \|\mathbf{B}\delta - \varepsilon\|_2^2 = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \varepsilon \in \mathbb{R}^{2N \times 1}. \quad (63)$$

3) *Temperature-dependent antenna delay variation*: Aside from the intrinsic variance in antenna delays across devices due to manufacturing, variations in antenna delays are influenced by temperature, with a gradient of approximately 7.17 ps/K [49]. On the lunar surface, temperatures undergo extreme variations, e.g., about 290 K at the equator between lunar noon and midnight [62]. This temperature swing can produce an antenna delay offset of up to 2.08 ns, corresponding to roughly 62.4 cm. However, this temperature variance occurs over 12 lunar hours, so the effective temperature change per Earth hour is smaller: 819 mK per Earth hour, which translates to a delay gradient of 5.87 ps per Earth hour (or approximately 1.76 mm per Earth hour). As a result, the theoretical upper bound for the required calibration interval can be determined based on the configured localization accuracy, Δd . For example, if $\Delta d = 1$ m, recalibration would be needed at most every 23.7 Earth days, equivalent to about 19.3 lunar hours.

D. Experiment

To calibrate the error introduced by antenna delays, it is necessary to first analyze its characteristics. Specifically, the distribution of the delay must be examined to determine its statistical behavior and consistency across measurements. In addition, its dependence on operational parameters should be systematically evaluated, including variations across pairwise combinations of nodes including distances, differences in message time intervals, and the influence of message scheduling strategies. Investigating these factors enables the identification of underlying patterns and dependencies, thereby providing the empirical foundation required to construct an accurate and reliable calibration model.

1) *Experimental Setup*: The experiment was conducted using four DWM1001-DEV nodes, labeled 1 through 4, arranged such that all nodes maintained unobstructed LOS with every other node. The inter-node distances were deliberately chosen to be distinct, with $d_{12} = 0.750$ m, $d_{13} = 0.550$ m, $d_{14} = 0.905$ m, $d_{23} = 0.620$ m, $d_{24} = 0.390$ m, and $d_{34} = 0.510$ m. All devices were connected to a centralized processing unit via wired UART communication at 115 200 baud, through




Fig. 11: Network topology of four nodes, each equipped with a DWM1000 UWB antenna, connected to a central processing unit via UART. Dashed and solid lines denote UWB and UART connections, respectively. Distances are not to scale.

which timestamp data were collected. The network topology employed in the experimental setup is illustrated in Figure 11.

Since the UWB transceiver on the DWM1001-DEV operates in half-duplex mode, each device was programmed with custom firmware that maintains a monotonically increasing clock derived from the DWM1001 tick counter, continuously monitors the channel for incoming messages, and transmits or resets its clock only upon explicit instruction from the centralized processing unit. The native tick counter is 40 bits wide, with each tick corresponding to 15.65 ps, yielding a maximum clock period of approximately 17.21 s before overflow. To accommodate longer measurement sessions, the firmware extends the counter to 64 bits, which expands the effective range to approximately 2.89×10^8 s.

When a DWM1001-DEV device transmits a message, the signal is broadcast to all compatible devices within its effective communication range. Since messages may be lost due to interference or timing uncertainties, each inter-node message is assigned a monotonically incrementing unique identifier, enabling the detection of packet loss and ensuring consistent temporal ordering of the collected timestamps. Each message additionally carries the identifier of the transmitting node to facilitate source attribution.

Each node reports a measurement to the centralised processing unit upon every transmission or reception event. A report comprises the node identifier (i.e., the source of the report), the transmitter address, the receiver address, and the message identifier. For transmitted messages, where no explicit receiver is designated, the transmitter's own address is recorded in the receiver address field. This convention ensures dataset consistency and provides a straightforward means of distinguishing transmitted from received messages within the collected records.

2) Results:

3) *Clock synchronization:* Prior to characterizing the systematic error, the clock synchronization performance among

the four devices is first examined. To ensure a controlled experimental setting, a round-robin message scheduling protocol is assumed, wherein each node transmits sequentially in ascending order of node index. It should be noted, however, that a node may occasionally transmit ahead of schedule due to computational throttling, introducing minor irregularities in the transmission order. (54) in online approach is additionally employed to track the evolution of the parameter estimates as the number of accumulated measurements grows.

Since the true clock parameters are not directly observable, where clock skew being governed by manufacturing tolerances and clock offset by systematic hardware conditions, their exact values cannot be known *a priori*. The estimation performance is therefore assessed indirectly through the internal consistency of the synchronization residuals rather than by comparison against ground truth.

Consequently, Figure 12 depicts the estimated clock parameters converging toward the true values for 30 ideal two-way measurements, in which each node transmits a message every 2 ms in a round-robin manner, rather than showing the true values directly. The final estimated clock parameters for Nodes 2–4 consist of clock skews equal to 1.00000678, 1.00000143, and 1.00000053, and clock offsets of -0.2593 s, -0.5593 s, and 11.24 s, respectively.

To further ensure robustness, the synchronization procedure is evaluated under varying message time intervals, denoted by T_k . This variation enables an assessment of the sensitivity of the clock parameter estimates to changes in the inter-message timing. Since clock offset is subject to systematic conditions, only the clock skews are observed. Figure 13 illustrates the final estimated clock skews as functions of the message time interval over 30 ideal two-way measurements averaged over 5 independent measurements. It is observed that the mean clock skews for Nodes 2–4 are 1.00000582, 1.00000086, and 1.00000121, respectively, with average standard deviations of 1.49×10^{-7} , 7.62×10^{-8} , and 1.09×10^{-7} over the four considered message intervals.

4) *Characterization of the systematic error: Multi-node analysis.* To characterize the systematic error, the distribution is derived through an offline estimation procedure where $K = 200$ two-way measurements between all unique pairwise node combinations are made with a time interval of 10 ms between messages, and node 1 is used as a reference node. To be consistent, a perfect and identical round-robin messaging protocol is reconsidered, wherein each node transmits sequentially according to its identifier. Using (54), the clock synchronization parameters are $\hat{\alpha}_{2-4} \approx 1.0$, $\hat{\beta}_2 = -0.258$ s, $\hat{\beta}_3 = -0.559$ s, and $\hat{\beta}_4 = 11.2$ s.

Figure 14 illustrates the residual distributions for all unique pairwise node combinations, with transmission directionality accounted for, and the corresponding expected values indicated by dashed red lines. The distributions likely attributable to unmodeled noise sources, and are therefore fitted using Gaussian Mixture Models (GMMs), shown as solid black lines. Examining the spread of these distributions, notable variation is observed across node pairs, in particular, pairs (1, 3), (1, 4), and (2, 4) exhibit standard deviations that exceed the target ranging accuracy of 1 m (equivalent to 3.34 ns).




Fig. 12: Iterative behavior of unconstrained clock synchronization, illustrating clock skews and offsets of nodes 2–4 with respect to node 1 over 30 ideal two-way measurements.




Fig. 13: Final estimated clock skews of Nodes 2–4 relative to Node 1 over 30 ideal two-way measurements for varying message intervals.

Distance analysis. To further characterize the systematic error and assess its dependence on range, node pair (1,2) was selected as a representative case. Figure 15 presents the evolution of the systematic error across varying separation distances. The residuals remain remarkably stable throughout, with the mean and standard deviation varying by no more than ± 35 ps and ± 51.9 ps, respectively. This corresponds to a range-equivalent error of $10.5 \text{ mm} \pm 15.6 \text{ mm}$. These results suggest that the systematic error can be reliably estimated at arbitrary distances, provided that the communicating nodes remain within each other’s effective communication radius.

Solving for the antenna delay per node using (56) and the means in Figure 14, we found the optimal TX and RX antenna delays for each node, which are presented in Tab. VI. The values found for the TX and RX antenna delays are similar to the results using the calibration algorithm in [49]. Using these values, the aggregated antenna delay terms in the measured timestamps as presented in (48a) and 48b.

To validate the clock synchronization and ranging performance, the optimal antenna delay compensation values were applied. The final clock parameter estimates and pairwise

TABLE VI: Estimated transmit (δ_T) and receive (δ_R) antenna delays per node using identical round robin message scheduling.

Node	δ_T (ns)	δ_R (ns)
Node 1	257.424321	257.444843
Node 2	257.087554	257.091556
Node 3	257.297166	257.291666
Node 4	257.434359	257.415335

distances were obtained using the proposed online–offline synchronization and ranging procedure with $K = 30$ two-way measurements. The estimated clock skews and offsets are given by $\hat{\alpha}_2 = 1.00000681$, $\hat{\alpha}_3 = 1.00000150$, $\hat{\alpha}_4 = 1.00000063$, $\hat{\beta}_2 = -0.2593$ s, $\hat{\beta}_3 = -0.5593$ s, $\hat{\beta}_4 = 11.24$ s. The corresponding estimated pairwise distances are $\hat{d}_{12} = 0.728$ m, $\hat{d}_{13} = 0.520$ m, $\hat{d}_{14} = 0.971$ m, $\hat{d}_{23} = 0.672$ m, $\hat{d}_{24} = 0.368$ m, $\hat{d}_{34} = 0.503$ m.

The estimated clock parameters are consistent with those reported in Sec. B.4.3. Furthermore, the estimated pairwise distances deviate from the ground-truth distances is at most 6.62 cm (between node 1 and node 4). This level of error re-




Fig. 14: Distributions of the residuals for all unique pairwise node combinations characterizing the systematic error from ingoing (blue) and outgoing (orange) measurements with clock synchronization parameters $\alpha_{2-4} \approx 1.0$, $\beta_2 = 16.95$ s, $\beta_3 = 16.64857729$ s, and $\beta_4 = 11.2$ s. Red dashed line denotes the expected value and solid black line denotes the fitted GMM distribution, per distribution.




Fig. 15: Distributions of the residuals between node 1 and 2 characterizing the systematic error from ingoing (blue) and outgoing (orange) measurements. Red dashed line denotes the expected value and solid black line denotes the fitted GMM distribution, per distribution.

mains within the accuracy requirements specified in Sec. B.2.1 despite the size of the residual distributions is larger than the required localization accuracy. This indicates that the proposed approach is suitable for calibration purposes under nominal operating conditions.

Robustness analysis. Nevertheless, these results primarily reflect performance under ideal conditions. A comprehensive evaluation must therefore also assess robustness against dis-




Fig. 16: Distributions of the residuals between node 1 and 2 characterizing the systematic error from ingoing (blue) and outgoing (orange) measurements using imbalanced round robin message scheduling where there are as half as many responses ($j \rightarrow i$). Red dashed line denotes the expected value and solid black line denotes the fitted GMM distribution, per distribution.

turbances and varying environmental conditions. One such scenario are missing messages due to NLOS conditions which causes multipath effects. This scenario is modeled as imbalanced round robin message scheduling where there are as half as many responses ($j \rightarrow i$). The message time interval is taken as 10 ms.

Figure 16 illustrates the characterization of the residuals using imbalanced round robin messaging schedule. The accuracy are different compared to Figure 14 where node pairs (1, 2), (2, 4), and (3, 4) exceeds the required localization accuracy. Moreover, the means are shifted as well compared to the previous result. However, the aggregated antenna delays over the ingoing and outgoing, $\Delta_{i \rightarrow j} + \Delta_{j \rightarrow i}$, means per node pair is identical to the previous result. This means that the TX and RX antenna delays can be identified and used to compensate the systematic error in the measurements. Verifying this, we found the estimated distances using the imbalanced round robin message scheduling as $\hat{d}_{12} = 0.723$, $\hat{d}_{13} = 0.512$, $\hat{d}_{14} = 0.954$, $\hat{d}_{23} = 0.678$, $\hat{d}_{24} = 0.365$, and $\hat{d}_{34} = 0.497$. The estimated distances have an error at most ± 5.81 cm which is well within the required localization accuracy.

E. Conclusion

UWB technology demonstrates considerable promise as a wireless communication solution for rover networks in space exploration scenarios. Nevertheless, inherent hardware imperfections introduce systematic antenna delays that must be rigorously characterized and compensated prior to reliable ranging operation. By formulating a centralized least-squares optimization that jointly performs clock synchronization and antenna delay estimation using a single reference node, the

proposed approach exploits a priori known pairwise distances to resolve the inherent identifiability ambiguity of the problem. Analysis of the resulting residual distributions yielded meaningful insight into the structure of the antenna delays and the characteristics of the underlying noise. The estimated antenna delays were subsequently applied to compensate the observed timestamps, yielding ranging accuracies of 6.62 cm and 5.81 cm under the ideal and imbalanced message protocols, respectively.

Future work should prioritize a systematic comparison with established state-of-the-art calibration methods, as well as experimental validation at inter-node distances more representative of operational lunar exploration scenarios. Furthermore, deploying the calibration procedure on resource-constrained embedded devices introduces additional practical challenges, particularly with respect to limited storage and computational capacity. In the present implementation, the transmission of raw timestamps to a centralized processing unit via UART introduces a communication bottleneck that adversely affects overall system efficiency. Addressing this limitation calls for a reformulation of the centralized calibration framework into a distributed and iterative architecture, wherein each node performs local computations while minimizing inter-node communication overhead, thereby enhancing scalability and real-time applicability in infrastructure-less environments.

APPENDIX C

PROOF OF STABILITY AND CONVERGENCE

This chapter presents a rigorous mathematical proof of the stability and asymptotic convergence of the proposed multi-agent coordination framework. The following proof evaluates the coupled interactions of the closed-loop system under the baseline dynamics, affine formation controls, auxiliary potential fields, and temporal adaptation laws derived in the main paper.

A. System Modeling and Global Error Space Formulation

Consider a multi-agent network consisting of N agents operating in a D -dimensional workspace. For a specific spatial coordinate component, let $\mathbf{z}_i \in \mathbb{R}^D$ and $\mathbf{v}_i = \dot{\mathbf{z}}_i \in \mathbb{R}^D$ denote the position and velocity vectors of agent i , respectively. Let $\mathbf{A}_z = K_d \mathbf{I}_D$ and $\mathbf{A}_z \mathbf{B}_z = K_p \mathbf{I}_D$, where $K_p, K_d > 0$ represent the baseline tracking gains.

By substituting the local temporal scaling factor $\tau_i(t) > 0$ and the transformation rules into the second-order acceleration dynamics given in (3) and (4), each tracking follower agent i yields the following unified model:

$$\tau_i \ddot{\mathbf{v}}_i = K_p (\mathbf{z}_i^{(T)} - \mathbf{z}_i) - K_d \mathbf{v}_i + \mathbf{f}_i(s_i) \quad (64)$$

$$+ \Phi_i^{\text{form}} + \Phi_i^{\text{obs}} + \Phi_i^{\text{agent}} + \Phi_{i,j}^{\text{conn}}, \quad (65)$$

$$\tau_i \dot{\mathbf{z}}_i = \mathbf{v}_i \implies \dot{\mathbf{z}}_i = \frac{1}{\tau_i} \mathbf{v}_i.$$

By setting the formation tuning matrices to match the baseline trajectory generator parameters ($\mathbf{A}_{\text{form}} = K_p \mathbf{I}_D$, $\mathbf{B}_{\text{form}} = K_d \mathbf{I}_D$ and fixing the normalization weight $\gamma_i = 1$, the formation term from (11) expands to:

$$\Phi_i^{\text{form}} = - \sum_{j \in \mathcal{N}_i} \omega_{i,j} (K_p (\mathbf{z}_i - \mathbf{z}_j) + K_d (\dot{\mathbf{z}}_i - \dot{\mathbf{z}}_j) - \ddot{\mathbf{z}}_j). \quad (66)$$

Let $\Omega \in \mathbb{R}^{N \times N}$ denote the signed stress Laplacian matrix characterizing the target affine configuration, which structurally satisfies $\Omega \mathbf{Z}^{(T)} = \mathbf{0}$. Let $\mathbf{s} = [s_1, \dots, s_N]^T \in \mathbb{R}^N$ be a collection of phase variables across the network. Collecting positions, velocities, accelerations, and learned forcing terms into global stacked vectors ($\mathbf{Z}, \dot{\mathbf{Z}}, \ddot{\mathbf{Z}}, \mathbf{F}(\mathbf{s}) \in \mathbb{R}^{DN \times 1}$) and assembling the local time-scaling terms into a diagonal matrix $\mathbf{T}(t) = \text{diag}(\tau_1 \mathbf{I}_D, \dots, \tau_N \mathbf{I}_D) \in \mathbb{R}^{DN \times DN}$, the global state acceleration space is expressed as:

$$\begin{aligned} \mathbf{T} \ddot{\mathbf{Z}} &= K_p (\mathbf{Z}^{(T)} - \mathbf{Z}) - K_d \mathbf{T} \dot{\mathbf{Z}} + \mathbf{F}(\mathbf{s}) \\ &\quad - (\Omega \otimes \mathbf{I}_D) (K_p \mathbf{Z} + K_d \mathbf{T} \dot{\mathbf{Z}} - \mathbf{T} \ddot{\mathbf{Z}}) \\ &\quad + \Phi^{\text{obs}} + \Phi^{\text{agent}} + \Phi^{\text{conn}}. \end{aligned} \quad (67)$$

We define the global target tracking deviation error as $\tilde{\mathbf{Z}} = \mathbf{Z} - \mathbf{Z}^{(T)}$, which dictates that $\dot{\tilde{\mathbf{Z}}} = \dot{\mathbf{Z}}$ and $\ddot{\tilde{\mathbf{Z}}} = \ddot{\mathbf{Z}}$. Substituting $\tilde{\mathbf{Z}}$ into (67) yields:

$$\begin{aligned} (\mathbf{I} - \Omega \otimes \mathbf{I}_D) \mathbf{T} \ddot{\tilde{\mathbf{Z}}} &= -K_p (\mathbf{I} + \Omega \otimes \mathbf{I}_D) \tilde{\mathbf{Z}} \\ &\quad - K_d (\mathbf{I} + \Omega \otimes \mathbf{I}_D) \mathbf{T} \dot{\tilde{\mathbf{Z}}} \\ &\quad + \mathbf{F}(\mathbf{s}) + \Phi^{\text{obs}} + \Phi^{\text{agent}} + \Phi^{\text{conn}}. \end{aligned} \quad (68)$$

Assuming the communication network is *affinely localizable* with a fixed configuration of leaders, the follower interaction submatrix is nonsingular, meaning the global transformation matrix inverse $\mathbf{M} = (\mathbf{I} - \Omega \otimes \mathbf{I}_D)^{-1}$ exists. Pre-multiplying (68) yields:

$$\begin{aligned} \mathbf{T} \ddot{\tilde{\mathbf{Z}}} &= -K_p \tilde{\mathbf{Z}} - K_d \mathbf{H} \mathbf{T} \dot{\tilde{\mathbf{Z}}} + \mathbf{M} \mathbf{F}(\mathbf{s}) \\ &\quad + \mathbf{M} \Phi^{\text{obs}} + \mathbf{M} \Phi^{\text{agent}} + \mathbf{M} \Phi^{\text{conn}}, \end{aligned} \quad (69)$$

where $\mathbf{H} = (\mathbf{I} - \Omega \otimes \mathbf{I}_D)^{-1} (\mathbf{I} + \Omega \otimes \mathbf{I}_D) \in \mathbb{R}^{DN \times DN}$ is a symmetric positive-definite matrix ($\mathbf{H} = \mathbf{H}^T > 0$).

B. Structural Properties of Spatial and Temporal Feedback Loops

Property 1: Velocity-Dependent Volumetric Dissipation. The volumetric obstacle repulsion Φ^{obs} (16) and mutual agent repulsion Φ^{agent} (17) are derived from velocity-dependent dynamic potential fields. Following their mathematical structure, their negative position gradients act as nonlinear damping elements that satisfy the inner product dissipativity conditions relative to physical velocities:

$$\dot{\mathbf{z}}^T \mathbf{M} \Phi^{\text{obs}} \leq 0 \quad \text{and} \quad \dot{\mathbf{z}}^T \mathbf{M} \Phi^{\text{agent}} \leq 0. \quad (70)$$

Furthermore, if an agent comes to rest ($\dot{\mathbf{z}} = \mathbf{0} \implies \mathbf{v}_i = \mathbf{0}$), these dynamic APF forces vanish identically: $\Phi_i^{\text{obs}} = \mathbf{0}$ and $\Phi_i^{\text{agent}} = \mathbf{0}$.

Property 2: Connectivity Graph Symmetry. Let the connectivity restoration gain matrix be $\mathbf{A}^{\text{conn}} = K_c \mathbf{I}_D$ with $K_c > 0$.

The stacked tracking feedback Φ^{conn} (19) is expressed using a symmetric, state-dependent connectivity Laplacian matrix $\mathbf{C}(\mathbf{Z}) = \mathbf{C}(\mathbf{Z})^\top \geq 0$. Because $\mathbf{C}(\mathbf{Z})\mathbf{Z}^{(T)} = \mathbf{0}$, it maps into error space as:

$$\Phi^{\text{conn}} = -(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D)\tilde{\mathbf{Z}}. \quad (71)$$

Property 3: Boundary and Monotonicity of Temporal Scaling. The proximity-based temporal adaptation law (20) dictates that if any agent i approaches an obstacle or a neighbor such that its minimum hazard clearance distance satisfies $d \leq d^{\text{slow}}$, its local time variable scales upward: $\tau^{\text{adapt}}(d) \geq \tau^{\text{base}} > 0$.

To ensure global network coordination during spatial deceleration maneuvers, agents run the discrete max-consensus protocol detailed in (23). Because $\tau^{\text{base}} > 0$ and $\tau^{\text{extra}} \geq 0$, the temporal matrix is strictly bounded and positive definite across all operation horizons: $\tau^{\min}\mathbf{I} \leq \mathbf{T}(t) \leq \tau^{\max}\mathbf{I}$, where $\tau^{\min} = \tau^{\text{base}}$.

C. Stability Analysis via the Lyapunov Direct Method

To establish the stability of the coupled spatial-temporal tracking states, we define a continuous energy-based Lyapunov scalar candidate function $V(\tilde{\mathbf{Z}}, \dot{\tilde{\mathbf{Z}}})$:

$$V = \frac{1}{2}K_p\tilde{\mathbf{Z}}^\top\mathbf{H}\tilde{\mathbf{Z}} + \frac{1}{2}\tilde{\mathbf{Z}}^\top(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D)\tilde{\mathbf{Z}} + \frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top\mathbf{T}(t)\dot{\tilde{\mathbf{Z}}}. \quad (72)$$

Because $\mathbf{H} > 0$, $\mathbf{C}(\mathbf{Z}) \geq 0$, and $\mathbf{T}(t) \geq \tau^{\min}\mathbf{I} > 0$, the function V is strictly positive definite for all nonzero tracking errors and satisfies $V(\mathbf{0}, \mathbf{0}) = 0$.

Differentiating V with respect to time along the continuous trajectories of our system yields:

$$\begin{aligned} \dot{V} &= K_p\tilde{\mathbf{Z}}^\top\mathbf{H}\dot{\tilde{\mathbf{Z}}} + \tilde{\mathbf{Z}}^\top(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D)\dot{\tilde{\mathbf{Z}}} \\ &\quad + \frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top(\dot{\mathbf{L}}_c(\mathbf{Z}) \otimes \mathbf{I}_D)\tilde{\mathbf{Z}} + \dot{\tilde{\mathbf{Z}}}^\top(\mathbf{T}\ddot{\tilde{\mathbf{Z}}}) + \frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top\dot{\mathbf{T}}\dot{\tilde{\mathbf{Z}}}. \end{aligned} \quad (73)$$

Substituting the decoupled state acceleration vector $\mathbf{T}\ddot{\tilde{\mathbf{Z}}}$ from (69) into (73) gives:

$$\begin{aligned} \dot{V} &= K_p\tilde{\mathbf{Z}}^\top\mathbf{H}\dot{\tilde{\mathbf{Z}}} + \tilde{\mathbf{Z}}^\top(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_d)\dot{\tilde{\mathbf{Z}}} \\ &\quad + \frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top(\dot{\mathbf{L}}_c(\mathbf{Z}) \otimes \mathbf{I}_d)\tilde{\mathbf{Z}} + \frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top\dot{\mathbf{T}}\dot{\tilde{\mathbf{Z}}} \\ &\quad + \dot{\tilde{\mathbf{Z}}}^\top \left[-K_p\mathbf{H}\tilde{\mathbf{Z}} - K_d\mathbf{H}\mathbf{T}\dot{\tilde{\mathbf{Z}}} \right. \\ &\quad \left. - \mathbf{M}(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_d)\tilde{\mathbf{Z}} + \mathbf{M}\mathbf{F}(\mathbf{s}) \right. \\ &\quad \left. + \mathbf{M}\Phi^{\text{obs}} + \mathbf{M}\Phi^{\text{agent}} \right] \end{aligned} \quad (74)$$

We simplify the terms using the algebraic structures of the unified feedback blocks:

- 1) By the symmetry of $\mathbf{H}$, the formation tracking potentials cancel out exactly: $K_p\tilde{\mathbf{Z}}^\top\mathbf{H}\dot{\tilde{\mathbf{Z}}} - K_p\dot{\tilde{\mathbf{Z}}}^\top\mathbf{H}\tilde{\mathbf{Z}} = 0$.
- 2) Because $\mathbf{M} = (\mathbf{I} - \Omega \otimes \mathbf{I}_D)^{-1}$ commutes with the tracking states and $\mathbf{C}(\mathbf{Z})$ is symmetric, the connectivity potential cross-terms cancel out identically within the network error space: $\tilde{\mathbf{Z}}^\top(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D)\dot{\tilde{\mathbf{Z}}} - \dot{\tilde{\mathbf{Z}}}^\top\mathbf{M}(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D)\tilde{\mathbf{Z}} = 0$.

- 3) The boundary transition matrix derivative satisfies $\frac{1}{2}\dot{\tilde{\mathbf{Z}}}^\top(\dot{\mathbf{L}}_c(\mathbf{Z}) \otimes \mathbf{I}_D)\tilde{\mathbf{Z}} \leq 0$, operating purely as a non-positive energy drop during network step-expansions.
- 4) During spatial adjustments near obstacles, the max-consensus discrete updates cause local temporal values to either remain constant or step upward to match a slower neighbor. Thus, during transitions, $\dot{\tau}_i(t) \geq 0$. To bound the transient impact of time-dilation variations on the tracking velocity layer, we apply a standard feedback gain padding constraint ensuring that $K_d\mathbf{H}\mathbf{T} - \frac{1}{2}\dot{\mathbf{T}} > 0$.

Applying these structural simplifications reduces the continuous energy derivative to:

$$\begin{aligned} \dot{V} &\leq -\dot{\tilde{\mathbf{Z}}}^\top \left(K_d\mathbf{H}\mathbf{T} - \frac{1}{2}\dot{\mathbf{T}} \right) \dot{\tilde{\mathbf{Z}}} \\ &\quad + \dot{\tilde{\mathbf{Z}}}^\top\mathbf{M}\mathbf{F}(\mathbf{s}) + \dot{\tilde{\mathbf{Z}}}^\top\mathbf{M}\Phi^{\text{obs}} + \dot{\tilde{\mathbf{Z}}}^\top\mathbf{M}\Phi^{\text{agent}}. \end{aligned} \quad (75)$$

1) Analysis of the Phase System Influence: To evaluate the driving DMP forcing vector $\mathbf{F}(\mathbf{s})$, we analyze the continuous canonical system across the network defined in (5). Let $\mathbf{s} = [s_1, \dots, s_N]^\top \in \mathbb{R}^N$. This can be written in global vector form using a positive semi-definite phase coupling Laplacian matrix $\Gamma \geq 0$:

$$\mathbf{T}_s\dot{\mathbf{s}} = -(\alpha_s\mathbf{I} + \Gamma)\mathbf{s}, \quad (76)$$

where $\mathbf{T}_s = \text{diag}(\tau_1, \dots, \tau_N) \in \mathbb{R}^{N \times N}$. We construct a local canonical phase validation function $V_s = \frac{1}{2}\mathbf{s}^\top\mathbf{T}_s\mathbf{s}$. Differentiating V_s gives:

$$\dot{V}_s = \mathbf{s}^\top\mathbf{T}_s\dot{\mathbf{s}} + \frac{1}{2}\mathbf{s}^\top\dot{\mathbf{T}}_s\mathbf{s} = -\mathbf{s}^\top \left(\alpha_s\mathbf{I} + \Gamma - \frac{1}{2}\dot{\mathbf{T}}_s \right) \mathbf{s}. \quad (77)$$

Because $\alpha_s > 0, \Gamma \geq 0$, and the maximum rate of change of temporal expansion $\dot{\tau}_i$ is naturally bounded by the smooth spatial approach profiles, the operational matrix $\left(\alpha_s\mathbf{I} + \Gamma - \frac{1}{2}\dot{\mathbf{T}}_s \right)$ remains strictly positive definite. Consequently, $\dot{V}_s < 0$ for all $\mathbf{s} \neq \mathbf{0}$, which mathematically guarantees that the coupled canonical states decay asymptotically to zero ($\lim_{t \rightarrow \infty} \mathbf{s}(t) = \mathbf{0}$). This ensures that the learned DMP spatial perturbations vanish entirely at steady state ($\lim_{t \rightarrow \infty} \mathbf{F}(\mathbf{s}) = \mathbf{0}$).

Substituting this result alongside Property 1 ($\dot{\tilde{\mathbf{Z}}}^\top\mathbf{M}\Phi \leq 0$) into (75) gives:

$$\dot{V} \leq -\dot{\tilde{\mathbf{Z}}}^\top \left(K_d\mathbf{H}\mathbf{T} - \frac{1}{2}\dot{\mathbf{T}} \right) \dot{\tilde{\mathbf{Z}}} \leq 0. \quad (78)$$

Since $\dot{V}$ is negative semi-definite, the integrated spatial-temporal system is *globally asymptotically stable in the sense of Lyapunov*. The tracking error states $\tilde{\mathbf{Z}}(t)$ and $\dot{\tilde{\mathbf{Z}}}(t)$ remain bounded ($\mathcal{L}_\infty$) for all $t \geq 0$.

D. Convergence and Safety Boundary Verification

1) Asymptotic Convergence to the Formation Configuration: Let $\Omega = \{(\tilde{\mathbf{Z}}, \dot{\tilde{\mathbf{Z}}} \mid \dot{V} = 0)\}$.

- 1) Given that $\left(K_d \mathbf{H} \mathbf{T} - \frac{1}{2} \dot{\mathbf{T}}\right) > 0$, the condition $\dot{V} = 0$ requires that the velocity tracking errors vanish identically: $\dot{\tilde{\mathbf{Z}}} \equiv \mathbf{0}$.
- 2) If $\tilde{\mathbf{Z}} = \mathbf{0}$ over an invariant time interval, the tracking acceleration vector must also vanish: $\ddot{\tilde{\mathbf{Z}}} \equiv \mathbf{0}$.
- 3) Once the transient phase passes, the canonical states decay ($\mathbf{F}(s) \rightarrow \mathbf{0}$). Away from obstacles, the APF vectors drop out ($\Phi^{\text{obs}}, \Phi^{\text{agent}} = \mathbf{0}$). Substituting these values back into the global error dynamics (69) yields:

$$\begin{aligned} \mathbf{0} &= -K_p \mathbf{H} \tilde{\mathbf{Z}} - \mathbf{M}(\mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D) \tilde{\mathbf{Z}} \\ &= -\mathbf{M}[K_p(\mathbf{I} + \mathbf{\Omega} \otimes \mathbf{I}_D) + \mathbf{C}(\mathbf{Z}) \otimes \mathbf{I}_D] \tilde{\mathbf{Z}}. \end{aligned} \quad (79)$$

- 4) Because $K_p > 0$, the matrix $(\mathbf{I} + \mathbf{\Omega} \otimes \mathbf{I}_D)$ is strictly positive definite via the affine localization requirement, and $\mathbf{C}(\mathbf{Z}) \geq 0$ is positive semi-definite. Their sum forms a strictly nonsingular, positive-definite operational matrix. Thus, the only mathematically valid state solution is:

$$\tilde{\mathbf{Z}} = \mathbf{0}. \quad (80)$$

By LaSalle's Invariance Principle, all tracking error states asymptotically converge to the origin. The network successfully achieves the target affine formation configuration at the baseline temporal rate:

$$\lim_{t \rightarrow \infty} \mathbf{z}_i(t) = \mathbf{z}_i^{(T)}, \quad \lim_{t \rightarrow \infty} \dot{\mathbf{z}}_i = \mathbf{0}, \quad \lim_{t \rightarrow \infty} \tau_i(t) = \tau^{\text{base}}. \quad (81)$$

2) Collision Avoidance and Local Minima Elimination:

If an agent approaches a volumetric boundary, the barrier parameter function goes to zero ($C(\Delta \mathbf{z}) \rightarrow 0$), driving the potential field to infinity ($\lim_{C \rightarrow 0} U_D = \infty$). Since $\dot{V} \leq 0$, the total closed-loop system energy is upper-bounded by its initial state ($V(t) \leq V(0) < \infty$). This restriction prevents the potential function from reaching infinity, mathematically guaranteeing that $C(\Delta \mathbf{z}) > 0$ for all $t > 0$, which proves that *collisions are completely avoided*.

Furthermore, if an agent enters a force-balance state near an obstacle and stops moving ($\dot{\mathbf{z}}_i = \mathbf{0}$), Property 1 dictates that the velocity-dependent avoidance forces completely deactivate ($\Phi_i^{\text{obs}} = \mathbf{0}, \Phi_i^{\text{agent}} = \mathbf{0}$). With the repulsive forces gone, the underlying DMP attractor and affine formation terms take back complete control of the acceleration layer ($\ddot{\tilde{\mathbf{Z}}} \neq \mathbf{0}$). This breaks the zero-velocity equilibrium, automatically allowing the agent to slide smoothly around the obstacles and *eliminating local minima trapping*. However, complete avoidance of local minima cannot be guaranteed in all scenarios. In particular, trapping may still occur when an agent is located on the far side of an obstacle along a line collinear with the goal direction, or when multiple obstacles form a non-convex configuration that generates a basin-like repulsive potential field. Under such geometric conditions, the combined obstacle-induced forces may still produce stable equilibrium that impede progress toward the goal.

3) *Connectivity Preservation Guarantee:* Suppose external avoidance maneuvers force two communicating agents apart such that $d_{i,j} > \alpha^{\text{conn}} r^c$. This activates the connectivity term $\Phi_{i,j}^{\text{conn}}$, adding a proportional restoring potential energy to V .

Because $V(t) \leq V(0) < \infty$, the maximum distance separation between any originally connected pair is structurally constrained by the initial energy allocation:

$$d_{i,j}(t) \leq \alpha^{\text{conn}} r^c + \sqrt{\frac{2V(0)}{K_c}}. \quad (82)$$

By choosing a sufficiently large connectivity restoration gain K_c such that $\sqrt{\frac{2V(0)}{K_c}} \leq (1 - \alpha^{\text{conn}}) r^c$, the maximum distance is strictly bounded:

$$d_{i,j}(t) < r^c \quad \forall t > 0. \quad (83)$$

This mathematically guarantees that the distance between neighboring agents never exceeds the maximum communication range r^c , proving that *global network connectivity is strictly preserved throughout all operations*.

To ensure system stability and asymptotic convergence during joint navigation, the effective connectivity activation threshold $\alpha^{\text{conn}} r^c$ must be strictly larger than the effective radius required for inter-agent collision avoidance. Accordingly, the following structural constraint must be satisfied:

$$\alpha^{\text{conn}} r^c \gg 2r^r + c^{\text{agent}}. \quad (84)$$


APPENDIX D SUPPLEMENTARY RESULTS AND FIGURES

A. Obstacle-Present Supervisory Demonstration

To evaluate the cooperative interaction between the AFC and the obstacle avoidance mechanism, the first experimental scenario was augmented by placing four volumetric obstacles around the two designated via-regions. These obstacles were positioned at coordinate pairs (0, 60) m, (30, 60) m, (100, 60) m, and (70, 60) m. The resulting system trajectories and corresponding tracking metrics are illustrated in Fig. 17.

Quantitative analysis indicates that despite local deflections caused by environmental hazards, the leader agents maintained a reference trajectory tracking error of 0.3442 m. Concurrently, the follower agents preserved the global geometric structure, restricting the affine formation tracking error to a sub-meter scale of 5.54 cm. Safety and network integrity were strictly maintained throughout the traversal. The minimum recorded obstacle clearance margin was 2.559 m, while the minimum communication distance margin remained bounded at 14.11 m, validating that the network successfully bypassed the newly introduced obstacles without fracturing the underlying communication topology.

The proximity of the obstacles to the target waypoints introduced a predictable tracking deviation, increasing the via-region clearance errors by 1.364 m for the waypoint at (20, 80) m and 2.230 m for the waypoint at (80, 80) m. Nevertheless, as demonstrated in Fig. 17, the multi-agent network successfully transited through both via-regions well within the specified mission tolerance radius of 10 m.



(a) Formation tracking state trajectories and chronological snapshots (a)-(f) outlining communication links, alongside time-history profiles along the x (g) and y (h) dimensions.




Fig. 18: Temporal of a hybrid-based robotic network navigation. Color map denote the selected network-wide temporal scaling. The temporal scalar laws are obstacle-based and inter-agent-based. Red denote the temporal scaling as function of the obstacles, and blue as function of the nearby agents.

time-scaling parameters, all uniformly assigned as $\tau_{\text{extra}}^{\text{agent}} = \tau_{\text{extra}}^{\text{obs}} = \tau_{\text{extra}}^{\text{form}} = \tau_{\text{extra}}^{\text{conn}} = 50$. However, in this demonstration, the obstacle and inter-agent temporal laws were visualized, where the corresponding the distance threshold are set to $d_{\text{slow}}^{\text{obs}} = 5$ m, and $d_{\text{slow}}^{\text{agent}} = 1$ m, respectively.

From Fig. 18, the obstacle-based temporal laws dominated the adaptive temporal synchronization since all agents were in close proximity with the obstacles which results in the agents to slow down around the cluttered environment. From Fig. 18(h), there is peak of inter-agent-based temporal scaling τ_{agent} around 10s, meaning that several agents were very close together. But τ_{obs} clearly dominates, which also ensures slowed movement around other agents.

After the network escaped the cluttered environment, as seen in Fig. 18(c)-(f), τ quickly reduces to the base time-scaling factor τ_{base} and remained in that value consistently, meaning that all agents were not too close to each other causing disturbance due to inter-agent collision avoidance and are merely converging slowly towards their respective target state. This demonstration that the adaptive temporal synchronization works as intended and enables safe traversal through cluttered environment and with dynamic agents.

C. Metric Distributions

Figures 19 through 22 illustrate the statistical distributions for each follower control configuration—specifically, *decentralized DMP*, *connected DMP*, *pure AFC*, and *hybrid DMP-AFC*—across all evaluation metrics in the local autonomy experiment, compiled over 500 independent evaluation runs per scenario.

B. Temporal Adaptation Demonstration

To demonstrate the working of the adaptive temporal synchronization τ , the robotic network navigated through a environment of Scenario C using the *Hybrid DMP-AFC* control law, as seen in Fig. 18. Similar to the local autonomy experiment, the base time-scaling factor of $\tau_{\text{base}} = 10$ was employed in conjunction with dedicated inter-agent, obstacle avoidance, formation tracking, and connectivity maintenance

(b) Robotic network performance metrics: (i) trajectory and formation tracking errors for leader agents (L_1 - L_3) and follower agents (F_4 - F_7), (j) minimum closest approach errors to via-regions verifying boundary penetration (green dashed line), (k) worst-case safety margins relative to the violation threshold (red dashed line), and (l) execution workloads per control loop.

Fig. 17: Comprehensive experimental results and performance evaluation of the robotic network during the first navigation task with obstacles around the two via-regions. The mission profile for leader agent L_1 originates at (20, 20) m, navigates through via-points at (80, 20) m and (80, 80) m, and terminates at (80, 10) m. Detailed formation, safety, and computational metrics distributions are tracked via horizontal bars (median), triangles (mean), and circles (outliers).




Fig. 19: Performance metrics across Scenarios A–C under Gain Configuration: Decentralized DMPs. Rows from top to bottom denote path length, formation error, communication, safety, and computation distributions.




Fig. 20: Performance metrics across Scenarios A–C under Gain Configuration: Connected DMPs. Rows from top to bottom denote path length, formation error, communication, safety, and computation distributions.




Fig. 21: Performance metrics across Scenarios A–C under Gain Configuration: Pure AFC. Rows from top to bottom denote path length, formation error, communication, safety, and computation distributions.




Fig. 22: Performance metrics across Scenarios A–C under Gain Configuration: Hybrid DMP-AFC. Rows from top to bottom denote path length, formation error, communication, safety, and computation distributions.