# Planning and Execution for Multi-robot Cyclic Rendezvous via Graphs of Convex Sets

Jaejin Cha and Dylan A. Shell

*Abstract*— **In multi-robot persistent monitoring problems, where each robot has a prescribed cyclic path, robots can improve overall performance by coordinating rendezvous communication while traversing their paths. Existing approaches based on consensus adjust all robots' traversals to obtain a single, common frequency. In this paper, we instead introduce a method based on having some robots take the specific role of accommodating others: this partitions robots into nominal robots that maintain a traversal frequency and courier robots that adjust their traversal frequencies to rendezvous with neighboring robots. Trajectories for couriers are planned in phase–time space by solving a shortest-path problem on a Graph of Convex Sets (GCS). Nominal robots' fixed frequencies promote stable, repeatable rendezvous patterns, while couriers provide flexibility needed to propagate information through the team. We evaluate the approach in simulation for monitoring performance, information sharing, and the time to reach an effective information-flow state. Experimental results show that role-based accommodation reaches an effective information-flow state quickly while maintaining high monitoring performance, with only minor degradation in information sharing. We also demonstrate the method on hardware using the Robotarium, showing that the planned phase-time trajectories can be executed on physical robots.**

## I. INTRODUCTION

In multi-robot persistent monitoring, a team of robots operates over long horizons in environments where events occur continuously and must be revisited regularly. Typical applications include surveillance [1], wildfire monitoring [2], and environmental monitoring [3] where performance is driven by repeated coverage and timely revisit rates. In many deployments, robots follow prescribed cyclic routes, for example pre-approved inspection loops, safe corridors, or patrol cycles, because fixed routes simplify safety guarantees, scheduling, and coverage analysis. When paths are fixed and cyclic, the remaining coordination variable is timing: traversal speed or equivalently traversal frequency along each cycle. This model reflects monitoring settings in which sensing requires physically visiting locations, and routes are selected to achieve the desired spatial coverage.

Communication is also fundamental for coordinated multi-robot systems. Monitoring performance, task allocation, and state estimation degrade when information becomes stale. In many settings, including underwater and underground environments, and disaster areas, communication is constrained by limited range or signal blockage. As a result, information exchange can be inseparable from motion. Robots communicate reliably only when they come within range, which motivates rendezvous-based communication in which robots exchange information within regions where communication is feasible. Under such constraints, robots should agree on a rendezvous plan at the point of last contact and commit to it as they cannot renegotiate meeting times while out of range. If a robot arrives at an agreed rendezvous while the other does not, the communication attempt fails. In persistent monitoring with cyclic paths, this requirement introduces a trade-off: rendezvous typically requires timing adjustments to meet neighbors, but these adjustments can degrade monitoring performance by distorting preferred cycle timing and revisit rates.

This paper examines this trade-off in a scenario where each robot repeatedly traverses a fixed cyclic path and communicates only within rendezvous regions induced by range limits. As the paths are fixed, each robot has a fixed set of potential communication partners determined by which paths come within communication range. The key control variable is each robot's traversal frequency around its given cyclic path. By adjusting traversal frequency (including varying it as a function of space or time), a robot can influence when it arrives at a rendezvous region to communicate, and therefore whether information propagates through the network.

An existing approach to maintain rendezvous communication under cyclic motion is to use consensus-based frequency synchronization. In [4], [5], when robots meet, they update their traversal frequencies using a local rule, and repeated interactions drive the network toward a shared frequency. While this can produce effective information propagation, enforcing a single common frequency can be undesirable for persistent monitoring. When paths are heterogeneous, a synchronized frequency may require some robots to move substantially faster or slower than their preferred monitoring rate, which distorts revisit times and can reduce monitoring performance. Moreover, synchronizing frequencies by convergence can take many rendezvous cycles, delaying the onset of when information is propagated reliably.

This paper proposes a new role-based accommodation approach that assigns communication responsibility to a subset of robots that are distributed in the network. Robots assigned the communication responsibility, named *couriers*, adjust their traversal frequency over time to accommodate rendezvous with their neighbors, while the complementary set, namely *nominal robots*, keep their traversal rate constant to preserve monitoring performance and ensure a predictable pattern. The couriers plan trajectories in phase-time space by solving a shortest-path problem on a Graph of Convex Sets (GCS) [6], producing repeatable rendezvous plans.

Address: Dept. of Computer Science & Engineering, Texas A&M University, College Station, TX, USA. {jaejin0109, dshell}@tamu.edu.

We evaluate the proposed method in simulation using metrics for monitoring performance, information sharing, and the time to reach an effective information-flow state. Results show that role-based accommodation reaches an effective information-flow state quickly while maintaining high monitoring performance, with only minor degradation in information sharing. We also demonstrate the method on hardware using the Robotarium platform [7], showing that the planned phase-time trajectories can be executed repeatedly on physical robots.

This paper's contributions are: 1.) Introducing a role-based formulation for cyclic rendezvous that partitions robots into roles, balancing monitoring performance and information propagation. 2.) Developing a planning method for courier accommodation via a GCS formulation to find a shortest-path [6], but in phase–time space, a representation different from how GCS was originally envisioned. 3.) Simulation and hardware results demonstrate that effective information flow is established quickly with minimal degradation in monitoring performance.

## II. RELATED WORK & BACKGROUND

Several papers motivate our focus on fixed cycles and timing control as primary design variables in persistent monitoring. Smith, Schwager and Rus [8] study persistent monitoring on a predefined closed path and control the robots' speed along a cycle that bounds an evolving region in space, and they optimize long-term monitoring performance under limited sensing. Employing multiple robots allows persistent monitoring workloads to split across agents. For example, balanced multi-UAV monitoring formulations employ explicit indicators and allocation procedures to reduce disparity in service difficulty and improve reliability over long missions [2]. Also, the closely related work of Alam and Bobadilla [9] also studies multi-robot coverage and persistent monitoring in sensing-constrained environments and proposes cyclic or periodic monitoring behaviors to repeatedly service regions of interest. These prior works, taken together, show how cyclic routes are natural in persistent tasks and this is a key modeling choice we also make. This leads directly to the question of how to influence task quality and burden distribution by optimizing timing along a cycle.

A common response to range-limited communication is to use intermittent rendezvous, in which robots execute their primary task while periodically meeting to exchange information. In [10], to enable communication-constrained exploration, they formalize the idea by alternating between exploration and rendezvous commitments using a Job-Shop Scheduling formulation to reduce waiting and improve information sharing without requiring full joint planning. Similarly, in [11], those authors emphasize rendezvous in communication-restricted, initially unknown environments. They design exploration behaviors which bias the team toward meeting sooner by leveraging properties specific to indoor spaces. Works such as these address the problem of *where* and *when* to meet while exploring an unknown environment.

Rendezvous planning also appears in persistent tasks, where meetings are needed not only to share information but to sustain operation with long duration. Mathew, Smith, and Waslander [12] address rendezvous planning for recharging in persistent surveillance, where a support team must meet working robots along their trajectories. Their formulation discretizes feasible rendezvous opportunities along known trajectories and solves a graph-based scheduling problem over a finite horizon.

In [4], [5], as noted earlier, the authors consider robots traversing predefined closed paths with communication possible only in rendezvous zones and use rendezvous-triggered speed synchronization to guarantee system-wide information propagation within a bounded time window. Their method provides a principled baseline for our local frequency synchronization methods and motivated our use of network-level delay metrics.

### A. Background: Graph of Convex Sets

The planning formulation we use is the GCS problem, introduced by Marcucci et al. [6]. In their framework, a planning problem is represented as a graph whose vertices are convex sets and whose edges encode feasible transitions between sets. The formulation allows each vertex and edge to be assigned both constraints and costs. Given some source and some destination, a shortest path algorithm on the GCS selects a sequence of vertices/edges and a specific point within each of the visited convex sets to minimize the total cost subject to feasible constraints.

## III. PROBLEM FORMULATION

We consider a team of robots executing fixed cyclic paths under a limited communication range, and we study how dynamically adjusting robots' traversal frequencies affects information propagation.

Let the robot team be indexed by $N = \{1, 2, ..., n\}$. Robots operate in world $W = \mathbb{R}^2$ or $W = \mathbb{R}^3$. Robot $i \in N$ is assigned a fixed cyclic path $\tau_i : S^1 \to W$, where $S^1$ denotes the unit circle $[0, 1)$. Cyclicity implies $\tau_i(0) = \lim_{x \to 1^-} \tau_i(x)$, and the robot's position in $W$ at time $t$ is given by $x_i(t) = \tau_i(\theta_i(t)) \in W$, where $\theta_i(t) \in S$ is the robot $i$'s phase on its cyclic path. As we have mentioned above, this formulation captures tasks where the path is fixed and the degree of freedom is how fast the robot traverses the cycle; it can be appropriate for problem instances in such applications as persistent environmental monitoring, perimeter patrol or industrial inspection.

Robots traverse along their cyclic path with a motion modeled by the phase dynamics $\dot{\theta}_i(t) = \omega_i(t)$, where $\omega_i(t)$ is the traversal frequency of robot $i$ at time $t$. In this work, the robots are assumed to be constrained to move only in the positive direction along the path parameterization (though this may be chosen to be either clockwise or anti-clockwise), so the traversal frequency is always non-negative and bounded $\omega_i(t) \in [\omega_i^{\min}, \omega_i^{\max}]$, where $\omega_i^{\min}, \omega_i^{\max} \in R_{>0}$ and $\omega_i^{\min} \leq \omega_i^{\max}$. These bounds are useful to encode both physical limits and monitoring requirements:

**Maximum frequency** captures physical limits of any robot. Let $v_i^{\max}$ be the speed limit of robot $i$, and let $L_i$ be the length of the fixed cyclic path. Then, $\omega_i^{\max} \leq \frac{v_i^{\max}}{L_i}$.

**Minimum frequency** represents the revisit constraint. Persistent monitoring often imposes a maximum allowable time between successive visits to points on the path. Let $T_i^{\max}$ be the maximum allowed lap time for robot $i$. Then, $\omega_i^{\min} \geq \frac{1}{T_i^{\max}}$.

Robots communicate using a range-limited wireless channel with communication radius $C > 0$. At time $t$, robots $i$ and $j$ can directly exchange information if their position is within the communication range $\|x_i(t) - x_j(t)\| \leq C$. Making the substitution $x_i(t) = \tau_i(\theta_i(t))$, this condition can be expressed in phase coordinates. The rendezvous region between robots $i$ and $j$ is defined to be the set of phase pairs for which communication is possible:

$$\Gamma_{ij} = \big\{(\theta_i, \theta_j) \in S \times S \big| \|\tau_i(\theta_i) - \tau_j(\theta_j)\| \leq C\big\}.$$

Because path $\tau_i$ and $\tau_j$ and communication range $C$ are fixed, $\Gamma_{ij}$ is a time-invariant set that can be computed offline from the environment and path geometry. Thus, here the rendezvous region is derived from physical range constraints and path geometry, which differs slightly from the formulation of [4] that inspired our work (they assume the rendezvous region is explicitly specified *a priori*). Although $\Gamma_{ij}$ is time-invariant, it is not assumed to be known initially. Robots discover their neighbors only through rendezvous events—once two robots meet, then they can use the observed path, phase, frequency, and role information in subsequent plans.

A rendezvous event occurs between robots $i$ and $j$ at time $t$ if $(\theta_i(t), \theta_j(t)) \in \Gamma_{ij}$. At such events, robot $i$ can observe the following for robot $j$: current phase $\theta_j(t)$ and frequency $\omega_j(t)$, and path $\tau_j$. For role assignment, we additionally assume robots can observe one another's current roles.

## IV. METHODS

In our treatment, all robots traverse their prescribed cyclic paths and continue to perform the monitoring task. Additionally, the method assigns one of two roles:

- **Nominal robots** traverse their cyclic paths at a fixed, default or preferred traversal frequency.
- **Courier robots** adjust their travel frequencies to ensure rendezvous with neighbors and thereby propagate information across teams.

Let $\rho_i(t) \in \{\text{nominal}, \text{courier}\}$ denote the role of robot $i$, and $\mathcal{N}_i(t)$ be the set of neighbors discovered by robot $i$ up to time $t$. Upon rendezvousing with robot $j$, robot $i$ records the timestamp $t$ and observes

$$\mathcal{O}_{ij}(t) = \big(\tau_j, \theta_j(t), \omega_j(t), \rho_j(t)\big).$$

The timestamped phase and frequency observation is used to predict the neighbor's future phase, while the observed role $\rho_j(t)$ is used by the local role-assignment rule described in Section IV-A.

Because nominal robots maintain a fixed traversal frequency, their timing provides repeatable rendezvous opportunities and is thus a stable reference against which couriers can plan. Courier robots vary their traversal frequencies to meet multiple neighbors and may therefore deviate from their preferred monitoring timing.

The proposed method consists of two parts. First, robots assign and update roles online at rendezvous events, with the goal of promoting interactions that are primarily between couriers and nominal robots (Section IV-A). Second, each courier computes a repeatable phase–time plan using GCS. The plan is reused until newly observed rendezvous information triggers replanning (Sections IV-B and IV-C).

### A. Role Assignment

The ideal local connectivity in the network is courier–nominal, where one robot actively accommodates rendezvous timing and the other maintains its preferred monitoring operation. Under communication constraints, roles can be assigned in an *ad hoc*, fully distributed manner based only on local rendezvous interactions.

When two nominal robots meet, the pair triggers a role-assignment step in which one robot becomes a courier. The choice can be made randomly (e.g., a coin flip) or via a heuristic that favors accommodation efficiency. The effectiveness of a courier assignment depends on factors such as the robot's part in the communication topology, the length of its path, and the frequency flexibility available for accommodation. If the communication graph and these costs were known in advance, the role assignment could be posed more explicitly as an allocation problem. In our setting, however, neighbor relations are *discovered* only through local rendezvous events, so robots must assign roles using information available at the moment of contact. Two useful heuristics we explored are: (i) assign the courier role to the robot with fewer neighbors (lower degree reduces accommodation burden), and (ii) assign the courier role to the robot with broader frequency bounds (greater flexibility improves its ability to accommodate rendezvous). Note that these heuristics employ information available locally at rendezvous. In the remainder of the paper, we report results employing only the simple, random assignment.

Whenever a courier meets a nominal robot, no role assignment is performed. The courier is responsible for accommodating rendezvous timing while the nominal maintains its fixed traversal frequency. If the courier has not previously considered that nominal neighbor in its accommodation plan, the courier replans to include the nominal in subsequent rendezvous.

When two couriers meet, we apply a special treatment to avoid increasing the dimensionality of coupling of the courier planning problem. Rather than forcing one courier to become nominal, each courier instead plans relative to a virtual nominal robot: an agreed reference schedule that does not correspond to a physical agent, but serves as a timing anchor that both couriers accommodate. The virtual nominal enables courier–courier rendezvous to be handled using the same courier–nominal planning interface, without requiring joint multi-courier optimization.

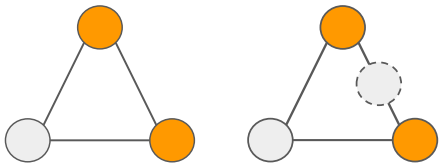

Fig. 1: A minimal case of non-bipartite communication network where a structure consisting only of courier–nominal connection is impossible. Orange vertices represent courier and gray vertex represents nominal. (a) Maintaining courier–courier connection. (b) Inserting virtual nominal between couriers. A dashed gray vertex represents virtual nominal.

The motivation for this construction is that the communication graph may not be bipartite, so it may be impossible to eliminate all same-role edges via role switching alone (see Figure 1). Bipartiteness is a global property of the graph, and the robots do not know the structure of the graph, so cannot fully verify this: they encounter and discover others only when traversing brings them physically within communication range. By inserting a virtual nominal on a courier–courier edge, we virtually enforce a courier–nominal interaction structure for planning purposes, allowing the network to be treated as consisting only of courier–nominal connections from the standpoint of accommodating planning, even when a true two-coloring does not exist.

### B. Expansion of Rendezvous Events

Now, we describe how courier robots accommodate their nominal neighbors. We begin with the rendezvous region induced by a pair of fixed cyclic paths and a communication range constraint, which is illustrated in Figure 2.

Given the rendezvous region $\Gamma_{ij}$ and the fixed traversal of a nominal robot $j$, we lift the feasible rendezvous opportunities from phase–phase space to phase–time space. Suppose that the most recent observation of robot $j$ was at time $t_0$. Its subsequent phase is

$$\theta_j(t) = \big(\theta_j(t_0) + \omega_j(t - t_0)\big) \bmod 1,$$

where $\omega_j$ is constant because robot $j$ is nominal. The corresponding phase–time rendezvous set for courier $i$ is

$$\mathcal{R}_{ij} = \left\{(\theta, t) \in S^1 \times [t_0, t_0 + T] \;\middle|\; \big(\theta, \theta_j(t)\big) \in \Gamma_{ij}\right\}.$$

Thus, both spatial compatibility and simultaneous arrival are enforced by requiring the courier trajectory to intersect $\mathcal{R}_{ij}$. This is illustrated visually in Figure 3.

The sets $\mathcal{R}_{ij}$ may be non-convex when the underlying paths overlap or intersect, as seen in Figure 4. For GCS planning, we decompose the rendezvous set into components and represent each component by a conservative maximum-area axis-aligned rectangle. This approximation may discard some feasible rendezvous opportunities, but it does not introduce rendezvous events that violate the original communication constraint.

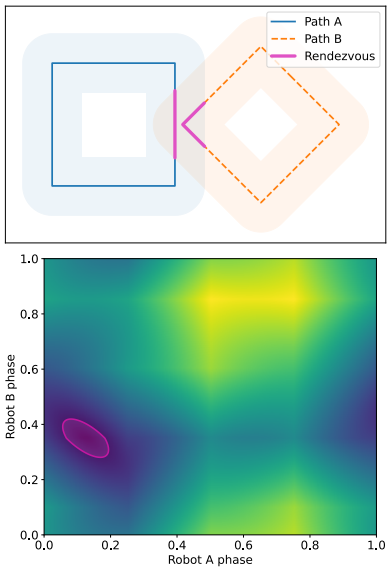


Fig. 2: An example of two cyclic paths with rendezvous region. (a) Two paths in world $W = \mathbb{R}^2$. (b) Euclidean distance plot in phase-phase space derived from the cyclic paths and communication range. Brightness shows the Euclidean distance $\|\tau_i(\theta_i) - \tau_j(\theta_j)\|$, and the magenta outline indicates the rendezvous region $\Gamma_{ij}$.

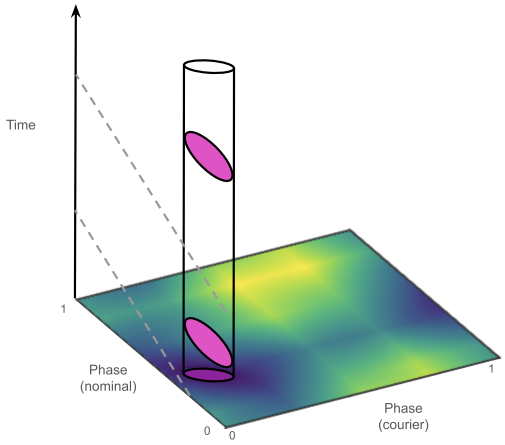


Fig. 3: Rendezvous events derived from a rendezvous region and a given nominal robot's traversal frequency, which can be visualized as straight lines because the nominal maintains a constant speed.

### C. Trajectory Planning via Graph of Convex Sets

The GCS framework has been used for planning trajectories in space, including in Euclidean space with dimensions greater than 2. We employ these techniques, but apply them to *spatio-temporal* coordination rather than purely geometric motion. Specifically, our planning is in a phase-time domain that captures the periodic nature of the rendezvous problem. This is inspired, in part, by the classic notion of a coordination space [13].

We now describe how to use GCS to plan courier trajectories in phase-time space. The GCS allows direct implementation for planning in Euclidean spaces, and rather than treating

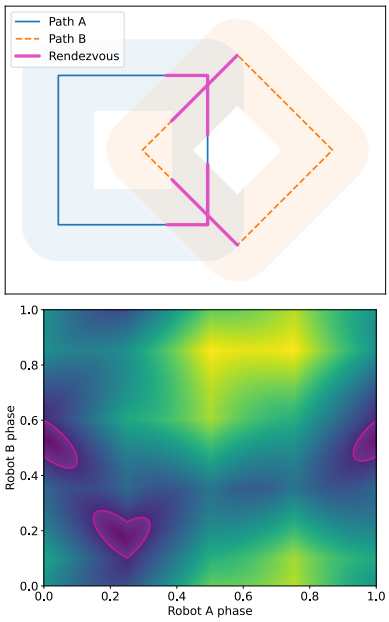


Fig. 4: An example of two cyclic paths that overlap. (a) Two paths in world $W = \mathbb{R}^2$. (b) Euclidean distance plot that contains non-convex rendezvous region $\Gamma_{ij}$.

$S^1$ as a manifold, we "unfold" a sufficiently large local region in order to plan cyclic paths. This is detailed next. We plan in unwrapped phase coordinates $s \in \mathbb{R}$: increasing $s$ by 1 corresponds to completing one lap and the wrapped phase is $s$ mod 1.

Because rendezvous opportunities repeat over time, we plan over a finite horizon and reuse the resulting plan periodically. We assume that the nominal-neighbor lap times repeat jointly after a finite interval; in our implementation, this is ensured by selecting traversal frequencies from a fixed-resolution grid. We choose this horizon as

$$T = \operatorname{lcm}\left\{ \frac{1}{\omega_j} \;\middle|\; j \in N_i^{\text{nom}} \right\},$$

where $N_i^{\text{nom}}$ denotes the nominal neighbors included in courier $i$'s current plan.

Over $[t, t+T]$ the courier plans a feasible phase-time trajectory that visits a subset of rendezvous events. At the end of the horizon, the courier transitions to a destination candidate that has the same wrapped phase as the start phase, allowing the plan to be repeated.

Given the horizon $T$, the courier's reachable phase range is bounded by its frequency limits. As the courier frequency satisfies $\omega \in [\omega^{\min}, \omega^{\max}]$, the unwrapped phase displacement satisfies $\Delta s \in [\omega^{\min} T, \omega^{\max} T]$. Hence, this bounds the phase-time domain that the courier can reach and therefore bounds the number of rendezvous events that need to be considered within the horizon.

With this limited phase-time space, rendezvous events become the vertices of the GCS. Each rendezvous vertex is represented as a convex region (or convex representation) described by inequality constraints. An example of the resulting planning space and rendezvous-event vertices is shown in Figure 5.

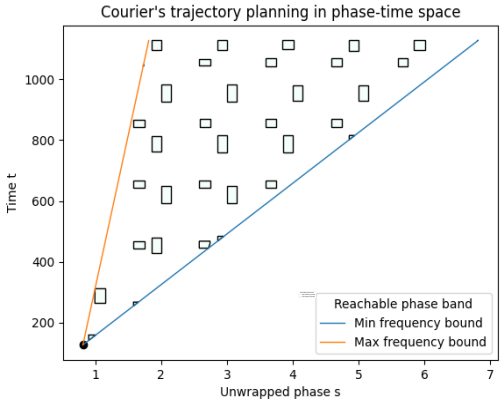


Fig. 5: The plots show the feasible planning space for the courier in phase-time (given $T$ and frequency bounds) and the available rendezvous events within that region.

Let the courier's current state in phase-time space be $(s, t)$. We construct a set of destination candidates at time $t + T$ whose wrapped phase equals the starting wrapped phase. In the unwrapped representation, this means destination of the form $(s+m, t+T)$ where $m$ is an integer chosen to respect the frequency bounds. From $\Delta s \in [\omega^{\min} T, \omega^{\max} T]$, the set of feasible integer displacements is

$$M = \left\{ \lceil \omega^{\min} T \rceil, \lceil \omega^{\min} T \rceil + 1, \ldots, \lfloor \omega^{\max} T \rfloor \right\}.$$

Therefore, the destination candidate set is

$$D(s,t) = \{(s+m, t+T) \mid m \in M\}.$$

We represent the source $(s, t)$ and each destination $(s + m, t+T)$ as singleton convex sets using equality constraints.

To run a shortest path algorithm, we construct directed edges between pairs of vertices (including source and destination). An edge from vertex $u$ at time $t_u$ to vertex $v$ at time $t_v$, is included only if it represents a physical feasible transition:

- **Forward-time constraint:** The edge must go to the future: $t_v > t_u$.
- **Frequency-bounded slope constraint:** The implied phase-time slope must be within frequency range:

$$\omega^{\min} \leq \frac{s_v - s_u}{t_v - t_u} \leq \omega^{\max}.$$

Costs can be assigned to both vertices and edges to modify courier behavior. We consider the following cost components:

- **Rendezvous count incentive:** Assign a constant negative cost to each edge (or to each rendezvous visit) so that the optimizer is encouraged to visit more rendezvous events within the horizon. The negative cost does not introduce negative cycle issues because the planning graph is acyclic, due to the forward-time constraint.

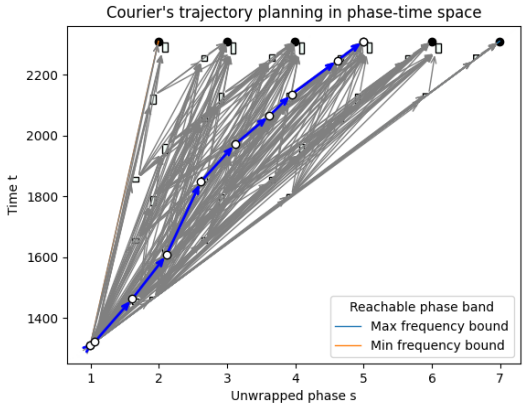


Fig. 6: Example GCS trajectory with two neighbors. The blue path represents an optimal plan produced by the solver.

- **Neighbor-frequency penalty:** Add a small penalty to rendezvous vertices associated with neighbors with high number of rendezvous events, encouraging visits to neighbors that are contacted less often and thereby reducing information staleness.
- **Spread cost:** Add a cost based on the squared phase displacement between consecutive vertices (an $L_2$ penalty) to discourage large jumps and promote even distribution in rendezvous selections.

With these costs and constraints, we solve a shortest path problem for each destination candidate in $D(s,t)$ and select the minimum-cost trajectory among them (see Figure 6).

### D. Techniques for Planning Efficiency

Under the explicit rendezvous-event representation, the size of the GCS graph becomes the main computational bottleneck. As the number of nominal neighbors grows, the planning horizon $T$ increases because it is chosen based on the repeat times of neighboring lap times. And, this increases in the number of rendezvous events, destination candidates, and feasible edges rapidly.

To manage computation time, graph-reduction techniques that may sacrifice some optimality for large gains in performance are worth considering. Perhaps the most direct way to reduce runtime is to sparsify the edge set:

- **Edge-length limit:** Restrict edges to connect only vertices whose time separation is below a threshold. Empirically, we observe that optimal trajectories often keep edge length short and visit rendezvous events relatively frequently, so very long edges are rarely needed.
- **Random edge removal:** Sample a subset of feasible edges uniformly (or alternatively via some heuristic priority) to reduce graph density.

Also, one can reduce the number of rendezvous vertices considered within the horizon:

- **Per-neighbor sub-sampling:** For each neighbor, retain the same fixed fraction of its rendezvous events. This preserves coverage across neighbors and reduces the risk that pruning eliminates rendezvous opportunities with a particular neighbor.

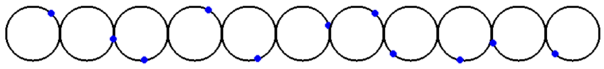

Fig. 7: A sparse, chain-like communication topology in which each robot has two neighbors except the endpoints. This map has 11 robots and stresses information propagation under limited graph connectivity.

- **Random sub-sampling:** Randomly discard rendezvous vertices to meet a computational budget.

Because both edge and vertex reduction can break connectivity between the source and destination candidates, we apply these reductions conservatively and verify that at least one feasible source-to-destination path remains. In the results we report, we employed edge-length limit and per-neighbor sub-sampling reduction techniques.

## V. SIMULATION EXPERIMENTS

### A. Implementation

We implemented a Python simulator that models robots traversing fixed cyclic paths, limited-range communication, and frequency updates at rendezvous events. The shortest path planning on a GCS used an existing Python library, GCSOPT [14]. We parameterize each robot's path as a Lissajous curve [15].

### B. Indicators of performance

We emphasize that GCS is only used for local courier planning, is not a centralized form of system-level optimization. As robots discover neighbors through range-limited rendezvous, we evaluate the distributed behavior that results using three indicators:

- **System information staleness:** The primary metric of communication performance. Intuitively, this measures how long it takes for information originating from one robot to become available to another robot through a sequence of rendezvous events. Let $I_{ij}(t)$ denote robot $i$'s most recent timestamp for information from robot $j$ at time $t$ (updated by exchanging timestamp records at rendezvous events). The system information staleness is defined as the maximum pairwise information staleness over all robot pairs:

$$IS_{\text{system}}(t) = \max_{i \neq j}(t - I_{ij}(t)).$$

- **Total frequency deviation:** We define the frequency deviation of robot $i$ at time $t$ as the absolute deviation from its initial frequency, $d_i(t) = |\omega_i(t) - \omega_i(0)|$.
  The total frequency deviation is the sum of per-robot deviations:

$$FD_{\text{total}}(t) = \sum_{i \in N} d_i(t) = \sum_{i \in N} |\omega_i(t) - \omega_i(0)|.$$

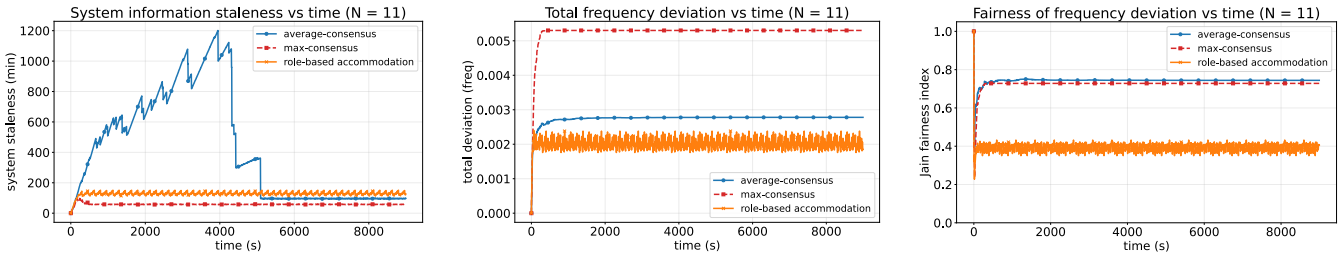


(a) System information staleness. (b) Total frequency deviation. (c) Jain's fairness index.

Fig. 8: Performance plots for the scenario shown in Figure 7 with 11 robots.

- **Fairness of frequency deviations:** To quantify how evenly deviation is distributed across robots, we use Jain's fairness index [16] to $d_i(t)_{i\in N}$:

$$J(t) = \frac{\left(\sum_{i=1}^{N} d_i(t)\right)^2}{N \sum_{i=1}^{N} d_i(t)^2},$$

where $J(t) \in (0, 1]$. We define $J(t) = 1$ when all deviations are equal, including the case where $d_i(t) = 0$ for all $i$. Larger values indicate that frequency deviation is more evenly distributed across the team.

### C. Results

We evaluated each method over 10 random-seed runs on the 11-robot sparse topology shown in Figure 7. The linear structure represents a worst-case problem for information propagation. The plotted time series report the mean across runs at each simulation time. For system information staleness, max-consensus uses all 10 runs, whereas average-consensus and role-based accommodation use nine successful runs; one run from each was excluded because frequency synchronization with a persistent phase offset prevented further rendezvous.

We compare with baseline methods (average-consensus and max-consensus) on the performance metrics described in the following. These baselines are:

1) **average-consensus:** Following [4], rendezvousing robots $i$ and $j$ set frequencies to their arithmetic mean:

$$\omega_i^+ = \omega_j^+ = \frac{\omega_i^- + \omega_j^-}{2}.$$

2) **max-consensus:** Rendezvousing robots $i$ and $j$ set both frequencies to their maximum:

$$\omega_i^+ = \omega_j^+ = \max(\omega_i^-, \omega_j^-).$$

The mean system information staleness is reported in Figure 8 (a). Max-consensus reaches an effective information-flow state fastest, closely followed by role-based accommodation, while average-consensus is slowest to do so. However, after reaching the effective state, role-based accommodation exhibits the highest information staleness among the compared methods.

The total frequency deviation metric across time appears in Figure 8 (b). Here, role-based accommodation yields the smallest overall deviation, indicating that the team remains closest to its monitoring-preferred timing when accommodation is concentrated on courier robots rather than distributed across all robots.

The fairness evaluation appears in Figure 8 (c). In this case, role-based accommodation achieves the lowest fairness, reflecting an intentionally uneven distribution of timing adjustment: couriers absorb most of the accommodation burden while nominal robots remain near their preferred frequencies.

## VI. DISCUSSION

We observe that the role-based accommodation method involves a trade-off. It reaches an effective information-flow state quickly, closely following max-consensus and considerably faster than average-consensus. Once this state is reached, however, it exhibits higher information staleness than the baseline methods. Still, the role-based method best preserves the robots' preferred monitoring frequencies.

The lower steady-state staleness of the consensus baselines is intuitive, since convergence to a common frequency forces rendezvous per cycle. However, average-consensus converges slowly, while max-consensus achieves rapid information flow at the cost of larger frequency deviation.

One can interpret the total frequency deviation as an indicator of overall monitoring performance: without any deviation at all, the robots are entirely ignoring the requirement to communicate and doing pure cyclic path execution. As the deviation increases, this represents the alteration of path timing in order to ensure rendezvous succeeds. The measure of fairness complements the deviation indicator by quantifying how equitable this alteration actually is. Thus, while role-based accommodation does attain good monitoring performance, this is at the cost of uneven compromise.

## VII. HARDWARE DEMONSTRATION

A hardware demonstration was conducted on the Robotarium [7]. Figure 9 validates that local role assignment and courier accommodation can be achieved without global frequency synchronization.

## VIII. CONCLUSIONS

This paper presented a role-based accommodation framework for multi-robot cyclic rendezvous under range-limited communication. By partitioning robots into some that maintain fixed traversal frequencies and others that adapt timing to meet neighbors, the method achieves good monitoring behavior while enabling rapid information propagation.

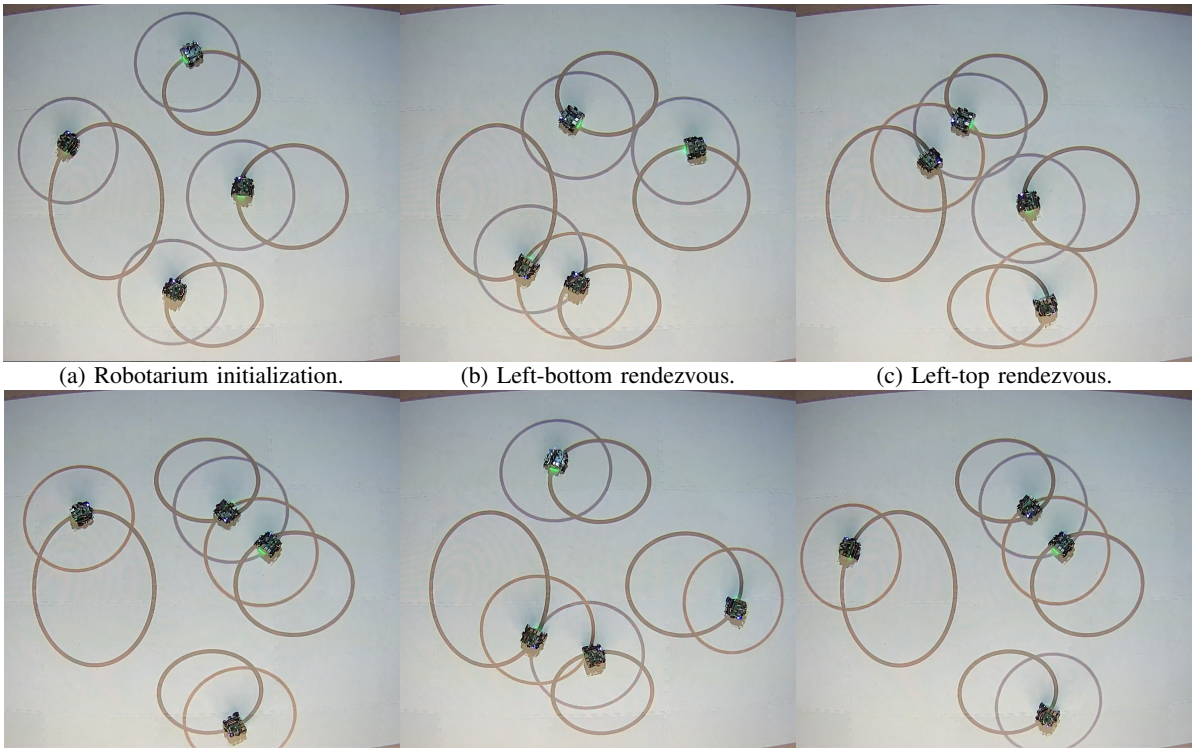

(a) Robotarium initialization. (b) Left-bottom rendezvous. (c) Left-top rendezvous.

(d) Top-right rendezvous. (e) Left-bottom rendezvous again (courier–courier). (f) Steady state.

Fig. 9: Four robots with a fixed neighbor structure in which each robot has two neighbors. In this demonstration, when courier–courier meet, they randomly reassign one to be a nominal, instead of creating virtual nominal.

Rendezvous strategies can be interpreted by how accommodation is distributed at meetings. Symmetric adjustments (e.g., average-consensus) can work when contacts are frequent, but in communication-limited settings they are often brittle because each robot's timing reference is changed by interactions with other robots, so commitments made at one rendezvous can be invalidated by subsequent interactions. This produces unstable plans and missed future meetings. Reliability improves when some robots remain consistent, while others absorb the timing effort.

## REFERENCES


[1] E. Stump and N. Michael, "Multi-robot persistent surveillance planning as a vehicle routing problem," in *IEEE International Conference on Automation Science and Engineering*, 2011, pp. 569–575.

[2] X. Zhan, Y. Chen, X. Chen, and W. Zhang, "Balanced multi-UAV path planning for persistent monitoring," *Robotica*, vol. 43, no. 1, pp. 332–349, 2025.

[3] J. Liu and R. K. Williams, "Monitoring over the long term: Intermittent deployment and sensing strategies for multi-robot teams," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7733–7739.

[4] X. Yu and M. A. Hsieh, "Synthesis of a time-varying communication network by robot teams with information propagation guarantees," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1413–1420, 2020.

[5] X. Yu, M. A. Hsieh, C. Wei, and H. G. Tanner, "Synchronous rendezvous for networks of marine robots in large scale ocean monitoring," *Frontiers in Robotics and AI*, vol. 6, p. 76, 2019.

[6] T. Marcucci, J. Umenberger, P. Parrilo, and R. Tedrake, "Shortest paths in graphs of convex sets," *SIAM Journal on Optimization*, vol. 34, no. 1, pp. 507–532, 2024.

[7] S. Wilson, P. Glotfelter, L. Wang, S. Mayya, G. Notomista, M. Mote, and M. Egerstedt, "The Robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems," *IEEE Control Systems Magazine*, vol. 40, no. 1, pp. 26–44, 2020.

[8] S. L. Smith, M. Schwager, and D. Rus, "Persistent monitoring of changing environments using a robot with limited range sensing," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 5448–5455.

[9] T. Alam and L. Bobadilla, "Multi-robot coverage and persistent monitoring in sensing-constrained environments," *Robotics*, vol. 9, no. 2, p. 47, 2020.

[10] A. R. Da Silva, L. Chaimowicz, T. C. Silva, and M. A. Hsieh, "Communication-constrained multi-robot exploration with intermittent rendezvous," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 3490–3497.

[11] M. Luperto, M. Tellaroli, M. Antonazzi, and N. Basilico, "Multi-robot rendezvous in communication-restricted unknown environments via backtracking and semantic frontier-based exploration," *Robotics and Autonomous Systems*, p. 105137, 2025.

[12] N. Mathew, S. L. Smith, and S. L. Waslander, "Multirobot rendezvous planning for recharging in persistent tasks," *IEEE Transactions on Robotics*, vol. 31, no. 1, pp. 128–142, 2015.

[13] P. A. O'Donnell and T. Lozano-Pérez, "Deadlock-free and collision-free coordination of two robot manipulators," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1989, pp. 484–489.

[14] T. Marcucci, "A unified and scalable method for optimization over graphs of convex sets," *arXiv preprint arXiv:2510.20184*, 2025.

[15] J.-A. Lissajous, *Mémoire sur l'étude optique des mouvements vibratoires*. Mallet-Bachelier, 1857.

[16] R. K. Jain, D.-M. W. Chiu, W. R. Hawe, *et al.*, "A quantitative measure of fairness and discrimination," *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA*, vol. 21, no. 1, pp. 2022–2023, 1984.