

Merry-Go-Round: Safe Control of Decentralized Multi-Robot Systems with Deadlock Prevention

Wonjong Lee¹, Joonyeol Sim², Joonkyung Kim², Siwon Jo³, Wenhao Luo³, and Changjoo Nam^{2,*}

Abstract—We propose a hybrid approach for decentralized multi-robot navigation that ensures both safety and deadlock prevention. Building on a standard control formulation, we add a lightweight deadlock prevention mechanism by forming temporary “roundabouts” (circular reference paths). Each robot relies only on local, peer-to-peer communication and a controller for base collision avoidance; a roundabout is generated or joined on demand to avert deadlocks. Robots in the roundabout travel in one direction until an escape condition is met, allowing them to return to goal-oriented motion. Unlike classical decentralized methods that lack explicit deadlock resolution, our roundabout maneuver ensures system-wide forward progress while preserving safety constraints. Extensive simulations and physical robot experiments show that our method consistently outperforms or matches the success and arrival rates of other decentralized control approaches, particularly in cluttered or high-density scenarios, all with minimal centralized coordination.

I. INTRODUCTION

As robots have been deployed to real-world environments, navigation of multiple robots has been an important and practical problem as its applications widely range from warehouse automation to search-and-rescue operations. Multi-Agent Pathfinding (MAPF) aims to generate globally optimal and deadlock-free paths for multiple agents/robots operating in a shared environment. Various centralized algorithms, such as Conflict-Based Search (CBS) [1] and its numerous enhancements (e.g., [2], [3]), offer solutions that explicitly resolve conflicts to ensure all robots reach their destinations without collisions or deadlocks.

Despite their effectiveness, conventional MAPF approaches face limitations that hinder deployment in real-world environments. Most of these methods rely on centralized computation and global replanning, requiring substantial communication bandwidth and computational resources. This becomes problematic when robots deviate from planned trajectories due to disturbances or obstacles. Moreover, these methods often neglect kinematic constraints, treating robots as holonomic point masses rather than considering nonholonomic motion. Scalability is another issue, as the state space grows exponentially with the number of robots, making real-time computation impractical. These challenges underscore the need for decentralized methods that distribute computation and enhance adaptability.

This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korea government (MSIT) (No. RS-2024-00461409 and No. 2022R1C1C1008476). ¹Dept. of Artificial Intelligence, Sogang University. ²Dept. of Electronic Engineering, Sogang University. ³Dept. of Computer Science, University of Illinois Chicago. *Corresponding author: c.jnam@sogang.ac.kr

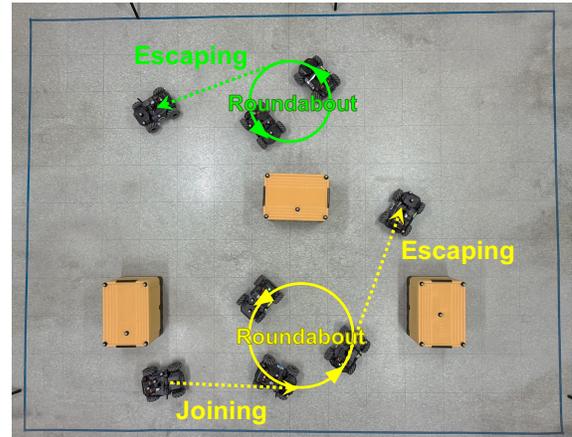


Fig. 1: An example where robots avoid deadlock situations using our proposed method — Merry-Go-Round. Robots encountering a potential deadlock form a roundabout and join it until they can escape to move ahead to their goal locations.

To address these limitations, decentralized navigation strategies, including variants of Artificial Potential Fields (APF) [4], Control Barrier Functions (CBF) [5], [6], and Optimal Reciprocal Collision Avoidance (ORCA) [7], [8], have gained attention for their ability to enable local decision-making without requiring a centralized planner. These methods allow robots to react to dynamic obstacles and nearby agents in real time to ensure safety. However, a key limitation of them is that they do not inherently prevent deadlocks but rather they happen to avoid them by ensuring local safety constraints. Recent works have discussed deadlock as undesired equilibrium [9], [10] and introduced centralized method for deadlock resolution [11]. Unlike centralized methods that explicitly enforce deadlock-free paths through global coordination, decentralized approaches lack a systematic mechanism to prevent deadlocks in cluttered environments.

We propose a hybrid approach that integrates decentralized reactive control with a higher-level deadlock prevention mechanism, aiming both safety and deadlock-free navigation. Our method operates by leveraging any control-based technique (e.g., [6]) to provide real-time collision avoidance while maintaining smooth motion. When a deadlock situation is predicted, a temporary roundabout maneuver is initiated among the involved robots, dynamically redirecting them until the situation is resolved. Fig. 1 illustrates some robots that join and escape their roundabouts. The roundabout maneuver prevents deadlocks by forming a dynamic circular path where robots rotate counter-clockwise while maintaining inter-robot safety constraints.

Our method offers several key contributions to multi-robot navigation: (i) it is fully decentralized, relying solely on

local peer-to-peer communication without requiring global coordination; (ii) it explicitly considers kinematic constraints, ensuring smooth motion for physical robots; (iii) it purposefully prevents deadlocks; (iv) its lightweight computation and local decision-making enable high scalability, making it effective in large-scale, dense environments; and (v) it is validated through extensive simulations and physical robot experiments, demonstrating its practicality and robustness.

II. RELATED WORK

MAPF has been studied as a centralized approach for computing optimal, deadlock-free paths. Algorithms like CBS [1] and its variants (e.g., ECBS [2], EECBS [3]) resolve conflicts via conflict trees and trajectory refinements. While ensuring completeness, they rely on a central planner, incurring high computational costs and requiring global communication. Thus, they exhibit limited scalability in dynamic or partially observable environments. Additionally, MAPF solution methods often ignore kinematic constraints, causing discrepancies in real-world deployments.

For better scalability and adaptability, decentralized approaches focus on local collision avoidance. APF [4] guides navigation with attractive and repulsive forces but struggles with local minima and oscillations. ORCA [7] extends velocity obstacles for real-time collision avoidance but lacks explicit deadlock prevention. CBFs [6] enforce safety by adjusting the controller of a robot via quadratic programming only when needed. PrSBC [5] enhances CBFs by handling measurement uncertainty, improving robustness in noisy environments. While these methods ensure safety, they lack explicit deadlock resolution or prevention.

Recent research has introduced learning-based approaches to enhance decentralized multi-robot navigation. Reinforced Potential Fields (RPF) [12] integrate reinforcement learning with APF. Graph Control Barrier Functions (GCBF+) [13] extend classical CBFs using graph neural networks. These approaches have demonstrated improved scalability and adaptability, particularly in highly dynamic settings. However, they still do not provide systematic deadlock prevention mechanisms, as their learned behaviors only mitigate deadlocks reactively rather than resolving them explicitly. Moreover, their generalization is limited, requiring retraining for unseen environments.

Some efforts have specifically targeted deadlock resolution in decentralized multi-robot navigation. Nonlinear Model Predictive Control (NMPC) [14] formulates deadlock prevention as a constrained optimization problem, ensuring that nonholonomic robots maintain collision-free paths while avoiding deadlocks. Unlike classical reactive methods, NMPC explicitly accounts for deadlock-prone scenarios by integrating a look-ahead optimization framework that anticipates future conflicts. However, the high computational complexity of NMPC significantly limits its scalability, especially when applied to large robot teams in real-time scenarios.

Despite advancements in decentralized safety and learning-based control, the intersection of (i) decentralized execution, (ii) kinematic feasibility, (iii) systematic deadlock

prevention, (iv) large-scale scalability, and (v) real-world deployment remains largely unexplored.

III. PROBLEM DESCRIPTION

A. System Model and Formulation

We consider N robots operating in a bounded 2D workspace $W \subseteq \mathbb{R}^2$. The set $O \subset W$ includes static obstacles. For each robot a_i where $i \in \{1, \dots, N\}$, let $x_i \in \mathbb{R}^n$ and $u_i \in \mathbb{R}^m$ be its state and control input, respectively. The system dynamics of a_i follow a control-affine form:

$$\dot{x}_i = f_i(x_i) + g_i(x_i)u_i \quad (1)$$

where $f_i : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g_i : \mathbb{R}^n \rightarrow \mathbb{R}^{(n \times m)}$ are locally Lipschitz continuous. The safety of system constraints is characterized through a safe set \mathcal{S} . For collision avoidance between robot pairs (a_i, a_j) where $i \neq j$, we define the pairwise safety function:

$$h(x_i, x_j) = \|p_i - p_j\|^2 - d_{\text{safe}}^2 \quad (2)$$

where $p_i \in \mathbb{R}^2$ represents the position components extracted from the full state x_i , which may include additional state variables such as orientation. A positive scalar value d_{safe} is the minimum allowed inter-robot distance. We set $d_{\text{safe}} = 2r_{\text{safe}}$ where r_{safe} is a safety margin for each robot. The global safety function is defined as:

$$\mathcal{S} = \{x_i \in \mathbb{R}^n \mid h(x_i, x_j) \geq 0, \forall j \neq i\}. \quad (3)$$

B. Control Lyapunov Functions

To stabilize the nonlinear system of robots, we employ Control Lyapunov Functions (CLFs). A continuously differentiable function $V : D \rightarrow \mathbb{R}_{\geq 0}$ is termed a CLF if it is positive definite and satisfies:

$$\inf_u [L_{f_i} V(x) + L_{g_i} V(x)u] \leq -\gamma(V(x)) \quad (4)$$

where $\gamma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a class \mathcal{K} function, meaning it satisfies $\gamma(0) = 0$ and γ is strictly increasing. This definition yields the following set of stabilizing controls:

$$K_{\text{clf}}(x) = \{u : L_{f_i} V(x) + L_{g_i} V(x)u \leq -\gamma(V(x))\}. \quad (5)$$

If V is a valid CLF, any locally Lipschitz continuous feedback $u = k(x)$ selected from $K_{\text{clf}}(x)$ will asymptotically stabilize the system to the equilibrium x^* such that $V(x^*) = 0$, typically at $x^* = 0$.

C. Control Barrier Functions

In the multi-robot system described above, stabilization focuses on driving the state of each robot to a desired configuration, while safety ensures that the system state remains within the safe set \mathcal{S} . A continuously differentiable function $h : X \rightarrow \mathbb{R}$ is termed a Control Barrier Function (CBF) if there exists an extended class \mathcal{K}_{∞} function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ (i.e., strictly increasing with $\alpha(0) = 0$) such that:

$$\begin{aligned} \sup_u [L_{f_i} h(x_i, x_j) + L_{g_i} h(x_i, x_j)u] \\ \geq -\alpha(h(x_i, x_j)), \forall j \neq i. \end{aligned} \quad (6)$$

The set of all safety-preserving controls can be expressed as:

$$K_{\text{cbf}}(x_i, x_j) = \{u : L_{f_i}h(x_i, x_j) + L_{g_i}h(x_i, x_j)u \geq -\alpha(h(x_i, x_j)), \forall j \neq i\}. \quad (7)$$

D. QP-based Control Formulation

To synthesize controllers that ensure both stability and safety, we formulate a Quadratic Program (QP) integrating the CLFs and CBFs, following the approach presented in [15]. Let $\mathbf{u} = (u, \delta) \in \mathbb{R}^m \times \mathbb{R}$ denote our decision variable where δ is a scalar relaxation variable. The QP formulation is

$$\min \frac{1}{2} \mathbf{u}^T H(x_i) \mathbf{u} + F(x_i)^T \mathbf{u} \quad (8)$$

$$\text{s.t. } L_{f_i}V(x_i) + L_{g_i}V(x_i)u + \gamma(V(x_i)) \leq \delta \quad (\text{CLF})$$

$$\begin{aligned} L_{f_i}h(x_i, x_j) + L_{g_i}h(x_i, x_j)u \\ + \alpha(h(x_i, x_j)) \geq 0, \forall j \neq i \end{aligned} \quad (\text{CBF})$$

where $H(x_i) \in \mathbb{R}^{(m+1) \times (m+1)}$ is a positive definite matrix ensuring the convexity of the QP, and $F(x_i) \in \mathbb{R}^{m+1}$ represents the linear term in the objective function. With δ , the stability constraint is allowed to be violated to maintain safety whenever necessary.

However, as noted in [10], even QP-based frameworks can admit undesirable equilibrium points that cause the system to remain stuck, resulting in a deadlock and preventing the robot from reaching its goal. To address this, we propose a deadlock prevention mechanism, which uses a roundabout maneuver, to escape such equilibria.

E. Assumptions

We assume all robots are homogeneous, sharing the same unicycle kinematics and communicating perfectly within δ_{comm} . Without loss of generality, their sensing range δ_{sensing} equals the communication range. Each robot knows the environment map, its own state, and its goal position. Owing to the barrier constraints, robots cannot approach obstacles within $2r_{\text{safe}}$. Consequently, goal positions must lie at least $2r_{\text{safe}}$ from obstacles, and likewise, any two goals must be separated by more than $2r_{\text{safe}}$, ensuring each pair of robots can converge to their respective goals without violating the barrier constraints. For handling obstacles, we use a right-hand rule which makes robots encountering static obstacles move clockwise around obstacles.

IV. MERRY-GO-ROUND FOR DEADLOCK PREVENTION

We propose the Merry-Go-Round (MGR) algorithm for multi-robot navigation while ensuring safety and preventing deadlocks. We first describe the algorithm and then provide a time complexity analysis.

A. Algorithm Description

The MGR algorithm consists of (i) deadlock prediction, (ii) avoidance circle (i.e., *roundabout*) generation, and (iii) escape condition checking that enables robots to safely exit the avoidance maneuver. The complete procedure is outlined in Alg. 1.

Algorithm 1 MERRY-GO-ROUND

Input: Robot a_i and its current position x_i , velocity v_i , and goal position g_i , a set of robots $a_j \in \mathcal{A}$ within δ_{comm} of a_i ($i \neq j$), time horizon T , the set of roundabouts \mathcal{C} , static obstacles O

- 1: **if** RECEIVEMGR(C) **then**
- 2: $a_i.\text{mode} \leftarrow \text{JOINMGR}(a_i, C)$
- 3: **else**
- 4: **for** each robot $a_j \in \mathcal{A}$ **do**
- 5: $\{x_j, v_j, g_j\} \leftarrow \text{RECEIVESTATE}(a_j)$
- 6: **if** ISDEADLOCKCANDIDATE(x_i, v_i, x_j, v_j, T) **then**
- 7: **if** ISGOALCHECKING($a_i, a_j, x_i, x_j, g_i, g_j$) **then**
- 8: **continue**
- 9: **end if**
- 10: $\mathbf{c} \leftarrow \text{FINDCENTER}(a_i, a_j)$
- 11: **if** $\exists C \in \mathcal{C}$ within δ_c of \mathbf{c} **then**
- 12: **if not** ISMGRVALID(C, O) **then**
- 13: $C \leftarrow \text{ADJUSTMGR}(C, O)$
- 14: **end if**
- 15: $a_i.\text{mode} \leftarrow \text{JOINMGR}(a_i, C)$
- 16: SENDMGR(a_j, C)
- 17: **else**
- 18: $C \leftarrow \text{CREATEMGR}(C)$
- 19: **if not** ISMGRVALID(C, O) **then**
- 20: $C \leftarrow \text{ADJUSTMGR}(C, O)$
- 21: **end if**
- 22: $a_i.\text{mode} \leftarrow \text{JOINMGR}(a_i, C)$
- 23: SENDMGR(a_j, C)
- 24: **end if**
- 25: **end if**
- 26: **end for**
- 27: **end if**
- 28: **if** ISESCAPABLE(a_i) **and** $a_i.\text{mode} = \text{MGR}$ **then**
- 29: $a_i.\text{mode} \leftarrow \text{ESCAPEMGR}(a_i)$
- 30: **end if**

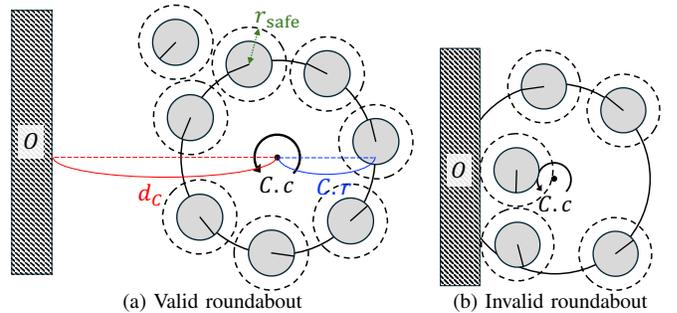


Fig. 2: Roundabouts for deadlock prevention. (a) Robots, each with a safe margin r_{safe} , form a roundabout. If the roundabout reaches capacity, additional robots orbit outside. The distance d_C ensures that all robots, including those on the outer orbits, remain at least $C.r$ away from the center $C.c$. (b) A roundabout can become invalid if d_C is too small, since the robots can no longer maintain the required distance $C.r$ from $C.c$.

Alg. 1 runs on a_i in a decentralized manner in each time a control input is calculated. It receives information regarding a_i (current position x_i , goal position g_i , and velocity v_i), nearby robots $a_j \in \mathcal{A}$, and other information (time horizon T for deadlock prediction, the set of roundabouts \mathcal{C} , and static obstacles O).

A roundabout C is a circular reference path for robots facing potential deadlock. Each robot travels around C in the same direction (counterclockwise in our implementation) until it reaches an escape condition. Ideally, each robot maintains a distance of $C.r$ from the roundabout center $C.c$. However, if C becomes full and cannot accommodate more

robots at radius $C.r$, additional robots orbit outside while still respecting the barrier constraints, as shown in Fig. 2a. Robots in outer orbits may later move into the inner orbit if there is sufficient space without violating the barrier constraints. As described in Fig. 2b, the robots may not be able to follow C if C is not sufficiently apart from O as the robots must ensure safety.

As a result of executing Alg. 1, each robot a_i sets its navigation mode $a_i.mode$ to either **GOAL** (moving directly to its goal) or **MGR** (following a roundabout). The controller then applies different strategies depending on the mode:

- **GOAL mode:** The nominal control input for a_i is computed based on feedback control using position error to guide the robot toward its goal position. The control law generates velocity commands in single integrator space, which are then converted to unicycle control inputs through state-dependent transformations in [16].
- **MGR mode:** The desired velocity \mathbf{v}_i^{des} for a_i is computed as the sum of a tangential velocity component \mathbf{v}_i^{tan} that drives counterclockwise circular motion and a radial velocity component \mathbf{v}_i^{rad} that maintains the desired radius from $C.c$:

$$\mathbf{v}_i^{des} = \frac{\mathbf{v}_i^{rad} + \mathbf{v}_i^{tan}}{\|\mathbf{v}_i^{rad} + \mathbf{v}_i^{tan}\|} v_{max} \quad (9)$$

where $\mathbf{v}_i^{tan} = v_{max} [-\sin \theta_i \cos \theta_i]$ with $\theta_i = \tan^{-1}[(y_i - c_y)/(x_i - c_x)]$ is the angle from $C.c$ to a_i . Additionally, $\mathbf{v}_i^{rad} = \frac{k_p}{C.n} (\|x_i - C.c\| - C.r) \frac{C.c - x_i}{\|C.c - x_i\|} v_{max}$ where $0 < k_p \leq 1$ is a proportional gain constant and $C.n$ is the number of robots current taking C .

In Alg. 1, a_i checks if any roundabout information C has been received from another robot (line 1), where C contains its center position $C.c$, radius $C.r$, and the number of robots currently taking C , indicated by $C.n$. If received, a_i immediately joins C (line 2) and the navigation mode of a_i remains **MGR** or switches to **GOAL** depending on the previous mode.

If no message has been received, Alg. 1 checks if there is a potential deadlock between a_i and a_j (line 6) according to two conditions:

- a) The distance between a_i and a_j is $2r_{safe}$ (e.g., Fig. 3a):

$$\|x_i - x_j\| = 2r_{safe} \quad (10)$$

which indicates that robots are already in barrier constraint active sets. In this case, the robots would be in an undesirable equilibrium, possibly causing a deadlock owing to their safety constraints being active, as proven in [10].

- b) The estimated trajectories of a_i and a_j are within $k_D r_{safe}$ over a time horizon T (e.g., Fig. 3b):

$$\min \|x'_i - x'_j\| \leq k_D r_{safe} \quad (11)$$

where $x'_i = x_i + v_i t_d$, $x'_j = x_j + v_j t_d$, and $0 \leq t_d \leq T$. The range of a constant k_D is $1 \leq k_D < 2$ because robots could enter the barrier constraint active sets

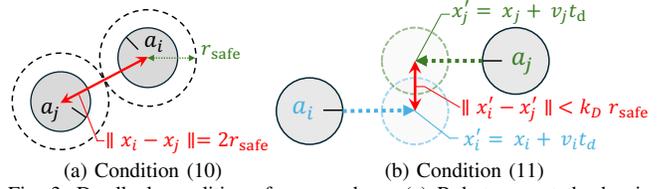


Fig. 3: Deadlock conditions for a_i and a_j . (a) Robots are at the barrier constraint distance $\|x_i - x_j\| = 2r_{safe}$. (b) There exists a time $t_d \in [0, T]$ where the estimated distance between a_i and a_j violates the barrier constraints.

of each other in this range, potentially leading to a deadlock situation.

In line 7, **ISGOALCHECKING** checks if a_i and a_j are already in close proximity (within $\epsilon < r_{safe}$) to their respective goal positions so they can reach their goals without a deadlock. If this condition is satisfied, the deadlock prevention for a_i and a_j is skipped. Specifically, the condition

$$\|x_i - g_i\| \leq \epsilon \text{ and } \|x_j - g_j\| \leq \epsilon \quad (12)$$

and $a_i.mode = a_j.mode = \text{GOAL}$

describes a situation where both robots are about to reach their destinations and not currently engaged in any roundabout maneuver.

If a potential deadlock is detected between a_i and a_j , **FINDCENTER** finds c which will be used as the initial center of a roundabout for the robots (line 10). We use the midpoint between a_i and a_j for simplicity, which can be calculated more accurately by predicting the intersection of the actual trajectories of the robots.

If there exists a roundabout C within proximity δ_c of c (line 11), a_i and a_j join C without generating a new roundabout. However, we must then verify that C is sufficiently distanced from O so that any robots using C can still navigate around $C.c$, even if C is already at capacity and additional robots must orbit on larger radii, as illustrated in Fig. 2a. In line 12, **ISMGRVALID** confirms $d_C \geq C.r + kC.n$, ensuring there is enough clearance from O to accommodate the outer orbits. If C is invalid, **ADJUSTMGR** finds a new center of C (line 13) such that $C.c$ and the closest obstacle in O is at least d_C . In our implementation, **ADJUSTMGR** discretizes a region of the environment around the current $C.c$ and searches for a valid location for a new $C.c$. It constructs a grid and marks each cell as either *valid* or *invalid*. A cell is marked *valid* only if it is at least $C.r + kC.n$ away from every obstacle. The valid cell with the lowest index among those having the closest distance to the current $C.c$ is selected as the new $C.c$.¹

Once a valid roundabout C is established, a_i joins C and its mode sets to **MGR** by **JOINMGR** (line 15). Since the radial component of (9) acts proportionally to the error between the current position and the target radius, the robot naturally moves along a spiral approach path when joining the roundabout, as depicted in Fig. 4. Simultaneously, it broadcasts C to robots within δ_{comm} (i.e., $a_j \in \mathcal{A}$) through **SENDMGR** (line 16), enabling a_j to join the same avoidance

¹An implementation of **ADJUSTMGR** can vary. A sampling-based method also can efficiently find a new $C.c$.

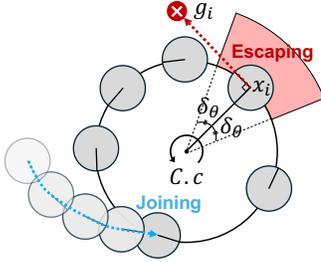


Fig. 4: The roundabout joining and escaping mechanism. Some robots joining the circular formation and others escaping when they meet escape conditions. The red shaded sector must be clear of other robots and obstacles for safe escape.

maneuver without requiring centralized coordination. If no nearby roundabout is found, a valid C is created (line 18). The minimum value of the radius $C.r$ is $2r_{\text{safe}}$ to ensure that the robots taking C do not collide with each other. A large value of $C.r$ is likely to lead to the generation of an invalid roundabout which overlaps with O .

In each control period, line 28 checks whether a robot a_i in MGR mode satisfies an escape condition, upon which a_i switches its navigation mode to GOAL (line 29). We formalize this condition based on the geometry among the roundabout center $C.c$, the robot position x_i , and its goal g_i . Let $\mathbf{v}_{ic} = C.c - x_i$ and $\mathbf{v}_{ig} = g_i - x_i$. If \mathbf{v}_{ic} is orthogonal to \mathbf{v}_{ig} , then a_i meets the geometric requirement for escape. However, a_i must also check the outer region (shown as the red shaded area in Fig. 4) is free from obstacles and other robots. This sector, centered at $C.c$, spans an angle $2\delta_\theta$ and extends to $\|C.c - x_i\| + \delta_{\text{sensing}}$. If this sector is free of obstacles or robots, a_i can safely escape C .

B. Time Complexity Analysis

We provide a proof for the time complexity of Alg. 1.

Theorem 4.1. For N robots, the computational complexity of Alg. 1 in each control period is $O(N^2 + N d_W d_H)$.²

Proof: Communication within δ_{comm} is assumed to be $O(1)$. We do not include the time for solving the QP, as our deadlock-prevention logic is layered on top of the controller.

In lines 4–26, the for loop iterates up to $N - 1$ times, once per neighboring robot. Within each iteration, ISGOALCHECKING checks up to N robots for their goal conditions and modes, contributing $O(N)$. Finding and creating a roundabout center via FINDCENTER and CREATMGR is constant-time, since these are closed-form operations. Checking validity with ISMGRVALID also involves a small, fixed set of geometric comparisons. Updating the navigation mode a robot in JOINMGR or ESCAPEMGR is $O(1)$.

In the worst case, ADJUSTMGR searches the entire discretized map of size $d_W \times d_H$, yielding $O(d_W d_H)$. Finally, ISESCAPABLE may need to check up to $N - 1$ robots to ensure no one blocks the outward path of escaping robot, adding $O(N)$.

Overall, we have $O(N(N + d_W d_H) + N) = O(N^2 + N d_W d_H)$. ■

²Here, d_W and d_H are the horizontal and vertical dimensions of the discretized map.

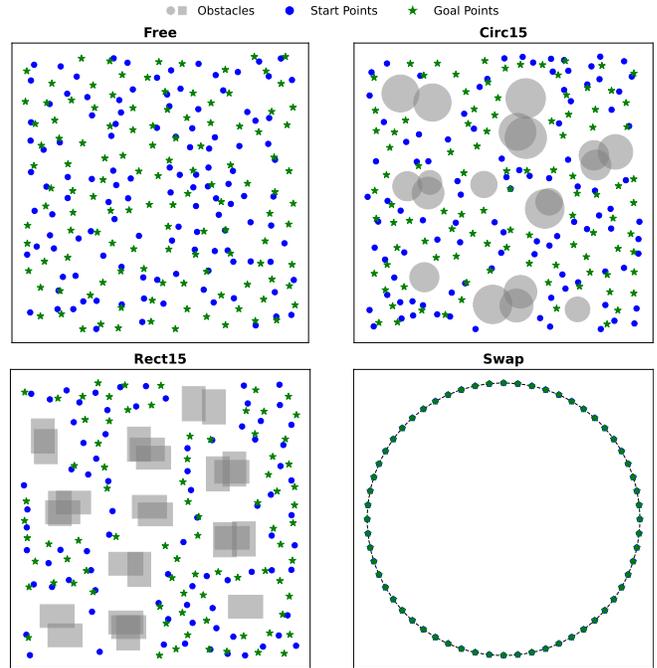


Fig. 5: The four environments for algorithm tests. (Clockwise from the top-left) They are Free (120 robots and their goals are plotted), Circ15 (100 robots), Swap (60 robots), and Rect15 (80 robots).

Practically, ADJUSTMGR does not need to scan the entire environment; in our experiments, searching a radius of $10C.r$ (i.e., about 3m in a $16\text{m} \times 16\text{m}$ map) around $C.c$ proved sufficient.

V. EXPERIMENTS

We evaluate the proposed algorithm in simulations and also with physical robots. In all experiments, algorithms run on a laptop with an AMD Ryzen 7 5800 8-Core Processor and 32GB RAM, using CVXOPT as the quadratic programming solver. The physical robot experiment is done with DJI RoboMaster S1 robots where their states are measured using a motion capture system (six NOKOV Pluto 1.3C).

A. Setup for Algorithm Tests

We construct a simulation environment where W is a $16\text{m} \times 16\text{m}$ area. Each robot has a radius of 0.2m, and the safe distance r_{safe} is 0.22m for enforcing barrier constraints. The maximum linear velocity of each robot is 0.8m/s, and its maximum angular velocity is $\pi/2$ rad/s. We test four different environments, as illustrated in Fig. 5. Free is an open space with no obstacles. Circ15 and Rect15 each contain obstacles that occupy 15% of W , circular in Circ15 and rectangular in Rect15. In these settings, the robot start and goal locations, as well as obstacles, are randomly generated for each instance. We generate 20 instances for each value of N (which is up to 120), with all methods evaluated on the same random instances for fair comparison. Swap employs the same obstacle-free environment as Free, but places the robots 15m from the center on opposite sides, requiring them to swap positions.

We compare our method with three widely used decentralized control techniques: GCBF+ [13], ORCA [7],

and CLF-CBF [15]. We use four metrics: (i) *success rate*, counting an instance as successful only if *all* robots reach their goals within a 2-minute limit; instances exceeding 2 minutes typically stall due to deadlock; (ii) *arrival rate*, the proportion of robots that reach their goals; (iii) *makespan*, the time until all robots have arrived in a successful instance (makespan is undefined in failed instances); and (iv) *mean time*, the average time among successfully arriving robots. Since non-arrivals are excluded, this metric can be skewed if many robots fail. If a method is with low arrival rates, the mean time only includes the result for relatively easy goals and simply discards challenging cases, thus artificially reducing their reported means.

For all experiments, we implement our QP formulation (8) with a positive definite Hessian matrix $H = \text{diag}(2, 2, 1)$. The linear term is chosen as $F = [-2(u_i^{\text{des}})^T, 0]^T$ for a_i . The CLF is implemented as a quadratic function $V(x_i) = (x_i - g_i)^T P(x_i - g_i)$ with $P = \text{diag}(1, 1)$, and $\gamma(V) = \lambda V$ where $\lambda = 1$. The CBF constraints are constructed using functions with $\alpha(h) = \beta h$ where $\beta = 5$. In addition, we use the following parameter values: deadlock prediction threshold $k_D = 1$, roundabout proximity $\delta_c = 2\text{m}$, roundabout radius $C.r = 0.3\text{m}$, communication range $\delta_{\text{comm}} = 1\text{m}$, radius increment constant $k = 0.1\text{m}$, and proportional gain $k_p = 0.05$. For the escape angle threshold δ_θ , we used different values depending on the environment: $\delta_\theta = \pi/6$ for environments with obstacles and $\delta_\theta = \pi/12$ for obstacle-free environments.

B. Results

Success and arrival results are summarized in Figs. 6 and 7 and Table I. Overall, our proposed method (MGR) consistently achieves superior or comparable performance across all environments and robot densities. In the obstacle-free setting (*Free*), MGR maintains near-perfect arrival and success rates even as the number of robots increases, outperforming CLF-CBF and GCBF+ and keeping pace with ORCA. In more difficult environments (*Circ15* and *Rect15*), MGR clearly outperforms the other methods: its arrival rates remain above 98%, and, crucially, it achieves a nonzero success rate for high robot counts—whereas most other methods drop to 0% success. Finally, in the *Swap* scenario, the performance of MGR remains robust, with success rates ranging from 95% to 100%, closely matching ORCA and surpassing CLF-CBF outright. This trend highlights the effectiveness MGR in both open and obstacle-dense environments, especially under high robot densities, where it maintains stable success and arrival rates that other methods struggle to achieve.

The makespan and mean navigation time are shown in Table II. MGR generally reports longer makespans and mean times — reflecting its more conservative roundabout-based approach — yet it consistently completes all scenarios, even with large teams or dense obstacles. In contrast, ORCA, GCBF+, and especially CLF-CBF often fail or yield infeasible times (shown as “-”) under higher robot counts or cluttered layouts. This trade-off demonstrates that the extra caution of MGR pays off in higher success rates and reliable

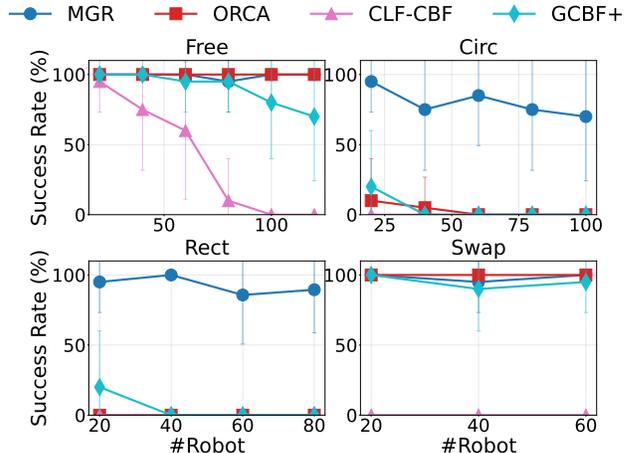


Fig. 6: The success rates of the compared methods

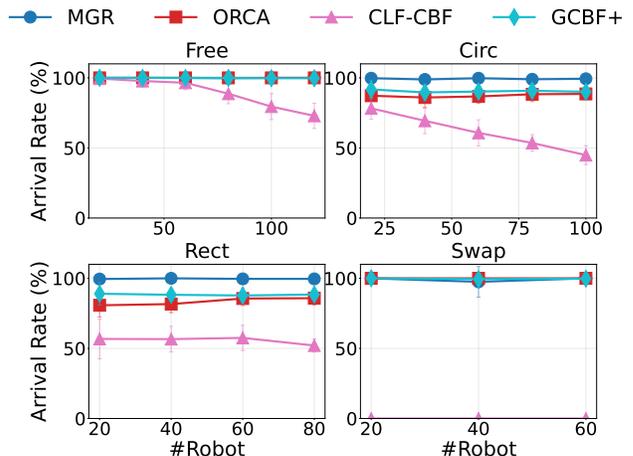


Fig. 7: The arrival rates of the compared methods

arrivals, making it more robust in challenging conditions where other methods struggle or stall.

We also attempted to compare with NMPC [14], which explicitly addresses deadlock prevention. However, NMPC becomes computationally prohibitive even at moderate team sizes: with just 20 robots, it requires over 10 seconds per robot for each control step, ruling out real-time implementation. In contrast, our approach remains efficient across all tests, never exceeding 2ms per robot per control period, even with 120 robots. This result indicates that our deadlock prevention strategy adds negligible overhead while still reliably avoiding deadlocks.

C. Physical Robot Experiment

For our physical robot experiments, we use a $3\text{m} \times 3.5\text{m}$ bounded workspace with robots (dimension is $320\text{mm} \times 240\text{mm}$). We run our method in three scenarios: *Free*, *Swap*, and *Rect*. The maximum linear velocity of each robot is 0.5m/s , and its maximum angular velocity is $\pi/2\text{rad/s}$. Although the physical robots support omnidirectional motion (including lateral movement), we restrict them to differential drive controls for consistency with our setup. A summary of the results can be found in the supplementary material.

VI. CONCLUSION

We proposed Merry-Go-Round (MGR), a decentralized multi-robot navigation framework that augments a standard

TABLE I: The success and arrival rates calculated from 20 instances for each team size and environment

Env.	#robot	Success Rate (%)				Arrival Rate (%)			
		MGR	ORCA	CLF-CBF	GCBF+	MGR	ORCA	CLF-CBF	GCBF+
Free	20	100.00	100.00	95.00	100.00	100.00	100.00	99.49	100.00
	40	100.00	100.00	75.00	100.00	100.00	100.00	97.75	100.00
	60	100.00	100.00	60.00	95.00	100.00	100.00	96.50	99.91
	80	95.00	100.00	10.00	95.00	99.75	100.00	88.68	99.93
	100	100.00	100.00	0.00	80.00	100.00	100.00	79.50	99.75
	120	100.00	100.00	0.00	70.00	100.00	100.00	72.91	99.75
Circ15	20	95.00	10.00	0.00	20.00	99.75	87.25	78.25	91.74
	40	75.00	5.00	0.00	0.00	98.87	86.00	69.37	89.62
	60	85.00	0.00	0.00	0.00	99.75	86.75	60.75	90.25
	80	75.00	0.00	0.00	0.00	99.00	88.31	53.56	90.81
	100	70.00	0.00	0.00	0.00	99.35	88.55	44.90	90.05
Rect15	20	95.00	0.00	0.00	20.00	99.49	80.75	56.75	89.00
	40	100.00	0.00	0.00	0.00	100.00	81.62	56.62	88.25
	60	85.00	0.00	0.00	0.00	99.60	85.58	57.50	87.75
	80	90.00	0.00	0.00	0.00	99.60	85.81	52.06	88.43
Swap	20	100.00	100.00	0.00	100.00	100.00	100.00	0.00	100.00
	40	95.00	100.00	0.00	90.00	97.50	100.00	0.00	99.49
	60	100.00	100.00	0.00	95.00	100.00	100.00	0.00	99.75

TABLE II: The makespan and mean time where the numbers in parenthesis are standard deviations

Env.	#robot	Makespan (sec)				Mean time (sec)			
		MGR	ORCA	CLF-CBF	GCBF+	MGR	ORCA	CLF-CBF	GCBF+
Free	20	21.19 (2.93)	18.38 (2.45)	20.32 (2.49)	31.37 (4.36)	12.79 (1.62)	9.97 (1.11)	11.59 (1.21)	16.53 (1.91)
	40	26.10 (3.01)	19.74 (1.88)	23.00 (2.36)	34.22 (3.40)	14.89 (1.30)	9.86 (0.73)	11.86 (0.78)	16.05 (1.28)
	60	32.15 (5.57)	22.31 (2.67)	33.69 (26.33)	36.48 (4.24)	18.04 (1.99)	10.74 (0.80)	13.10 (1.18)	16.86 (0.93)
	80	44.17 (11.03)	22.93 (2.07)	28.05 (0.30)	36.59 (2.79)	22.07 (1.57)	11.01 (0.55)	13.38 (1.23)	17.13 (1.08)
	100	53.35 (12.15)	25.73 (4.19)	-	39.77 (3.45)	27.70 (2.80)	11.80 (0.79)	14.05 (1.37)	17.51 (0.87)
	120	60.44 (10.85)	28.98 (3.09)	-	40.81 (3.67)	33.29 (3.21)	12.65 (0.85)	14.54 (1.19)	17.76 (0.79)
Circ15	20	31.68 (7.06)	49.90 (0.85)	-	33.93 (3.81)	17.35 (3.10)	13.14 (2.64)	12.38 (2.40)	18.47 (3.75)
	40	39.93 (8.05)	100.60 (0.00)	-	-	21.02 (2.65)	15.14 (2.63)	12.19 (1.26)	18.38 (1.56)
	60	53.22 (7.88)	-	-	-	27.54 (2.50)	15.09 (1.71)	12.73 (1.71)	19.52 (1.37)
	80	67.76 (12.69)	-	-	-	33.14 (3.39)	17.49 (2.38)	13.06 (1.90)	18.80 (1.22)
Rect15	20	37.63 (11.00)	-	-	41.87 (8.25)	19.15 (3.44)	13.23 (2.69)	11.59 (1.12)	19.09 (2.60)
	40	46.30 (8.81)	-	-	-	23.99 (2.86)	13.65 (1.58)	12.30 (2.06)	19.44 (2.14)
	60	57.17 (10.44)	-	-	-	31.00 (4.46)	15.47 (2.18)	13.02 (1.67)	19.21 (1.50)
	80	69.90 (13.48)	-	-	-	39.11 (4.19)	15.50 (1.68)	12.86 (1.74)	19.15 (1.17)
Swap	20	37.67 (4.45)	32.15 (0.00)	-	36.84 (0.00)	30.38 (2.10)	26.93 (0.00)	-	33.72 (0.03)
	40	44.77 (10.12)	42.35 (0.00)	-	41.8 (5.59)	33.96 (6.36)	34.82 (0.00)	-	35.63 (0.95)
	60	52.61 (13.31)	45.10 (0.00)	-	44.39 (3.32)	36.22 (3.39)	35.63 (0.00)	-	36.52 (0.36)

controller with explicit deadlock prevention. By detecting deadlocks, robots form local roundabouts and move in a controlled orbit until safe to proceed. This lightweight, peer-to-peer approach is highly scalable. Extensive simulations and physical robot tests show that MGR consistently achieves high success and arrival rates, outperforming or equaling existing decentralized controllers, especially in dense or cluttered environments where deadlocks commonly occur.

REFERENCES

- [1] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [2] M. Barer, G. Sharon, R. Stern, and A. Felner, "Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem," in *Proc. of Int. Symp. on Combinatorial Search*, vol. 5, 2014, pp. 19–27.
- [3] J. Li, W. Ruml, and S. Koenig, "EECBS: A bounded-suboptimal search for multi-agent path finding," in *Proc. of AAAI Conf. on Artificial Intelligence*, vol. 35, 2021, pp. 12 353–12 362.
- [4] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. of Robotics Research*, vol. 5, pp. 90–98, 1986.
- [5] W. Luo, W. Sun, and A. Kapoor, "Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates," in *Proc. of Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 372–383.
- [6] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Trans. on Robotics*, vol. 33, pp. 661–674, 2017.
- [7] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Proc. of Int. Symp. on Robotics Research*, 2011, pp. 3–19.
- [8] D. Bareiss and J. Van den Berg, "Generalized reciprocal collision avoidance," *Int. J. of Robotics Research*, vol. 34, pp. 1501–1514, 2015.
- [9] J. S. Grover, C. Liu, and K. Sycara, "Deadlock analysis and resolution for multi-robot systems," in *Proc. of Workshop on the Algorithmic Foundations of Robotics*, 2021.
- [10] M. F. Reis, A. P. Aguiar, and P. Tabuada, "Control barrier function-based quadratic programs introduce undesirable asymptotically stable equilibria," *IEEE Control Systems Letters*, vol. 5, pp. 731–736, 2020.
- [11] M. Jankovic, M. Santillo, and Y. Wang, "Multiagent systems with cbf-based controllers: Collision avoidance and liveness from instability," *IEEE Trans. on Control Systems Technology*, vol. 32, pp. 705–712, 2023.
- [12] D. Zhang, X. Zhang, Z. Zhang, B. Zhu, and Q. Zhang, "Reinforced potential field for multi-robot motion planning in cluttered environments," in *Proc. of IEEE/RSSJ Int. Conf. on Intelligent Robots and Systems*, 2023, pp. 699–704.
- [13] S. Zhang, O. So, K. Garg, and C. Fan, "GCBF+: A neural graph control barrier function framework for distributed safe multi-agent control," *IEEE Trans. on Robotics*, 2025.
- [14] A. S. Lafmejani and S. Berman, "Nonlinear mpc for collision-free and deadlock-free navigation of multiple nonholonomic mobile robots," *Robotics and Autonomous Systems*, vol. 141, p. 103774, 2021.
- [15] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. on Automatic Control*, vol. 62, pp. 3861–3876, 2017.
- [16] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The robotarium: A remotely accessible swarm robotics research testbed," in *Proc. of IEEE Int. Conf. on Robotics and Automation*, 2017, pp. 1699–1706.