ICRA2026 Late Breaking Results Poster presented at
2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Lateral Reciprocal Collision Avoidance: A Probabilistic Social-Norm-Inspired Strategy for Deadlock-Free Multi-Robot Navigation

Siyi Lu[1], Sipu Ruan[*1]

I. Introduction and Motivation

Safe and efficient multi-robot navigation is fundamental in robotics, essential for many service and industrial applications. In such tasks, multiple robots share an environment and must reach their goals without colliding with each other. Designing strategies that guarantee collision-free and efficient motion remains challenging. Multi-robot navigation methods are either centralized or decentralized. Centralized approaches offer strong safety and optimality but are computationally expensive, rely on synchronized communication, and scale poorly. Decentralized methods use only local information (e.g., positions, velocities), making them lightweight, robust, scalable, and more suitable for real-world applications.

Among decentralized approaches, the Velocity Obstacle (VO) framework is widely used for collision detection between robot pairs. Building on VO, Reciprocal Collision Avoidance (RCA) methods share avoidance responsibility, enabling efficient and scalable navigation. However, RCA-based methods suffer from a critical limitation—deadlock—which can occur even between two robots, severely degrading success rates and efficiency. From an optimization perspective, RCA methods can be seen as VO-based quadratic programs (QPs), where collision avoidance is a hard constraint and goal reaching is the objective. This safety-first formulation guarantees collision-free behavior but can lead to deadlock—robots remain safe but make no progress. In this paper, we analyze deadlock in RCA frameworks through VO-based QPs and provide a theoretical characterization using the KKT conditions.

Unlike robotic systems, human crowds exhibit robust collision avoidance and rarely experience deadlock, even in dense settings. Pedestrian behavior is inherently decentralized, relying on local observations and shared conventions—i.e., social norms. A key assumption underlying these norms is that *all agents follow the same avoidance strategy*. One particularly effective norm is symmetric lateral displacement, where individuals consistently sidestep in a coordinated manner, which plays a crucial role in preventing deadlock. Motivated by this observation, we propose incorporating pedestrian social norms into the RCA framework to mitigate deadlock in multi-robot navigation. Specifically, as shown in Fig. 1, we introduce Lateral Reciprocal Collision Avoidance (LRCA), a novel RCA-based strategy that embeds symmetric lateral displacement as a shared consensus among robots. By integrating this social norm into VO-based decision-making, LRCA significantly reduces deadlock while preserving decentralization. Moreover, we provide a theoretical analysis based on VO-based QPs and KKT conditions to demonstrate that LRCA effectively avoids deadlock.

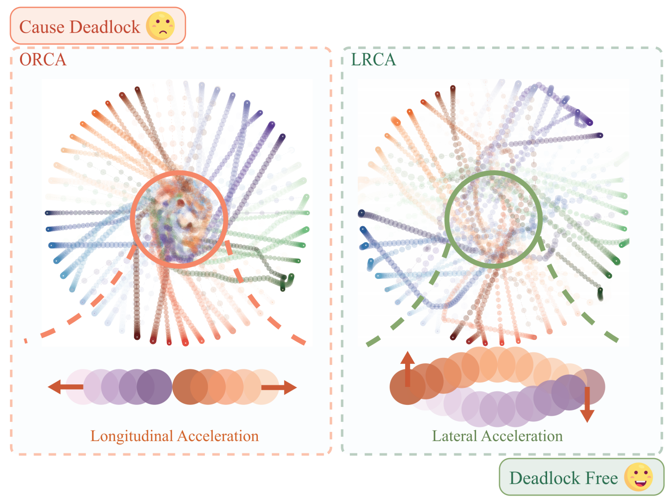

Fig. 1. Left: 50 robots navigating with ORCA, which relies solely on longitudinal acceleration and is prone to deadlock. Right: 50 robots with LRCA, which incorporates lateral acceleration to mitigate deadlock.

- We formalize the RCA framework as a VO-based QPs and provide a theoretical analysis of deadlock scenarios using KKT conditions.
- We propose LRCA, a novel RCA framework inspired by pedestrian lateral avoidance behavior.
- We prove and validate LRCA's ability to avoid deadlock in several scenarios prone to ORCA deadlock.

II. Methodology

In the proposed framework, we aim to define sets of permitted velocities V_i for robot i and V_j for robot j, such that these sets ensure reciprocal collision avoidance within a given time step τ.

[1]Siyi Lu, Sipu Ruan are with the Robotics Institute, School of Mechanical Engineering and Automation, Beihang University, China. siyilu@buaa.edu.cn, ruansp@buaa.edu.cn.
[*]Sipu Ruan is the corresponding author.

A. Two Random Lateral Avoidance Strategies

Inspired by pedestrian collision avoidance behavior, we constrain the relative velocity adjustment $\delta \boldsymbol{v}$ to avoid being parallel to the current relative velocity, thereby preventing deadlock.To prevent the difference in velocity adjustments between two robots $\delta \boldsymbol{v}_i - \delta \boldsymbol{v}_j$ from being parallel to the relative velocity $\boldsymbol{v}_i^{opt} - \boldsymbol{v}_j^{opt}$, two cases are considered: both robots' velocity adjustments can be in the clockwise direction of the relative velocity, or both can be in the counterclockwise direction, as shown in that part of the picture in the poster.

Since LRCA introduces directional constraints on velocity adjustments (clockwise or counterclockwise relative to the relative velocity), selecting the minimum adjustment to avoid collisions would shrink the feasible velocity set and may even lead to an empty set, causing another form of deadlock. To address this, we instead select the adjustment with the largest angle relative to the relative velocity—specifically, the velocity perpendicular to the linear boundary of the $VO^{\tau}_{A|B}$. This maximizes the feasible set while ensuring collision avoidance. As illustrated in the poster, the adjustment can be perpendicular to either the left or right linear boundary of $VO^{\tau}_{i|j}$, corresponding to two deadlock avoidance strategies.

To enable robots to reach consensus without explicit communication, selecting either Case 1 or Case 2 , we introduce a randomized mechanism inspired by human game-theoretic interactions. Each robot i randomly selects whether the velocity adjustment $\delta \boldsymbol{v}_i$ lies on the clockwise or counterclockwise side of the relative velocity with respect to robot j. If both robots i and robot j choose the same side, collision avoidance is achieved without deadlock; otherwise, avoidance is not possible.

B. Theoretical Analysis of Deadlock

Using the QP formulation and KKT conditions, we prove that LRCA is deadlock-free for two-robot interactions. The key insight is that when the preferred velocity $\hat{v}_i \neq 0$, the stationarity condition yields:

$$v_i^{new*} = \hat{v}_i + \frac{1}{2}\lambda_{ij}^* n$$

where n is the chosen lateral normal vector. Because n is perpendicular to the relative position vector (rotated by $\frac{\pi}{2} \pm \theta$), it has a non-zero component orthogonal to the line connecting the robots. Consequently, the optimal velocity v_i^{new*} cannot be zero unless $\hat{v}_i = 0$. This contrasts with ORCA, where the normal aligns with the relative position, permitting a zero-velocity solution in symmetric deadlock configurations.

III. Simulation Results

We evaluated LRCA against four baselines (ORCA, T-MPC, SARL, RL-RVO) in three benchmark scenarios: Cross (symmetric intersection), Corridor (narrow passage), and Random (dense, unstructured). Key results are summarized below.

A. Effectively Avoid Deadlock

In center-symmetrical crossing scenarios, ORCA causes deadlock: each robot moves straight toward its goal, and ORCA enforces symmetric deceleration to avoid collision. This causes both robots to gradually slow down and stop at the safety boundary with zero relative velocity, preventing them from reaching their goals despite nonzero target velocities. In contrast, the proposed LRCA breaks symmetry via randomized lateral avoidance, enabling smooth passing without stopping. By restricting velocity changes to a randomly selected side of the relative velocity, LRCA induces early asymmetric behavior, thus avoiding convergence to the zero-velocity equilibrium that causes deadlock.

B. Navigation Tasks for Different Scenarios

The picture in this part of the poster presents qualitative comparisons of navigation trajectories generated by ORCA, T-MPC, SARL, RL-RVO, and the proposed LRCA in three representative scenarios. In the cross scenario (first row), LRCA produces smooth and well-structured trajectories by introducing consistent lateral avoidance, enabling robots to pass through the intersection efficiently. In the corridor scenario (second row), LRCA maintains clear directional separation, allowing robots to traverse the corridor smoothly with minimal interference. In the random scenario (third row), LRCA achieves compact and well-separated trajectories, demonstrating robust collision avoidance without sacrificing efficiency.

C. Excellent Navigation Performance

In the cross scenario, LRCA consistently outperforms other methods across all metrics. The success rate for LRCA nearly 100% for all group sizes, while ORCA sees a significant decline, dropping from 100% to 70.30% as the number of robots increases. For extra distance, LRCA shows the best performance, with values ranging from 0.26 meters for 4 robots to 1.37 meters for 50 robots. Regarding extra time, LRCA again leads, with values ranging from 2.91 seconds for 4 robots to 6.27 seconds for 50 robots. Finally, in terms of average speed, LRCA maintains high speeds, from 0.8 m/s for 4 robots to 0.58 m/s for 50 robots.

D. High Computational Efficiency

The picture in this part of the poster shows the computation time (mean/standard deviation in milliseconds) for different methods with varying numbers of agents. ORCA is the most computationally efficient method, with times increasing from 1.69×10^{-5}ms for 2 agents to 1.76×10^{-3}ms for 60 agents. T-MPC has a higher computational cost, increasing from 7.01×10^{-3}ms to 1.55×10^{-1}ms as the number of agents grows. SARL exhibits the highest computational cost, rising from 1.39×10^{-1}ms for 2 agents to 8.75 seconds for 60 agents. RL-RVO has moderate computation times, ranging from 3.72×10^{-3}ms to 1.21×10^{-1}ms. Our method, LRCA, provides a good balance, with times increasing from 3.26×10^{-5}ms to 1.21×10^{-2}ms, making it efficient yet slightly slower than ORCA.