IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Velocity Potential Field Modulation for Dense Coordination of Polytopic Swarms and Its Application to Assistive Robotic Furniture

Lixuan Tang, David Rüegg, Runze Zhang, Anastasia Bolotnikova, Jan Rabaey, Auke Ijspeert

***Abstract*—We explore the use of a mobile furniture swarm that are intended to assist users with limited mobility in their daily indoor activities. We focus on the multi-robot coordination problem when a dense target pose configuration is required, such as in an apartment setting. In those cases, the convergence of one robot to the target can be significantly affected by neighboring robots with specific shapes. In this letter, we propose a solution, named Velocity Potential Field Modulation (VPFM), to deal with the dense coordination problem of a polytopic swarm in a decentralized manner. We adapt our method to assistive applications, such as room reconfigurations and facilitating indoor movement of wheelchair users. We evaluate the performance of our method in simulations and on real-world mobile furniture hardware, demonstrating its effectiveness and real-time performance.**

***Index Terms*—Collision Avoidance, Path Planning for Multiple Mobile Robots or Agents, Swarm Robotics**

I. Introduction

DAILY life presents numerous challenges for the elderly and individuals with disabilities, as limited mobility affects many aspects of their experience. For instance, wheelchair users often struggle to navigate and avoid obstacles in cluttered indoor environments filled with space-occupying furniture. It is also difficult for them to organize the room by manual effort. Providing mobility to indoor furniture is a promising way to ease their life. Instead of introducing new assistive agents, we aim to integrate assistive functionalities into everyday objects already present in indoor environments. As illustrated in Fig. 1, mobile furniture pieces can function as fundamental components of an interactive assistive environment, adapting to occupants' needs and preferences. To facilitate multiple uses of limited space, mobile furniture pieces can autonomously rearrange themselves from one formation to another (e.g., setups for meetings, parties, meals, or indoor cleaning). To enhance daily autonomy, assistive furniture can actively move out of the way for a wheelchair user passing by, or follow the user to help carrying objects.

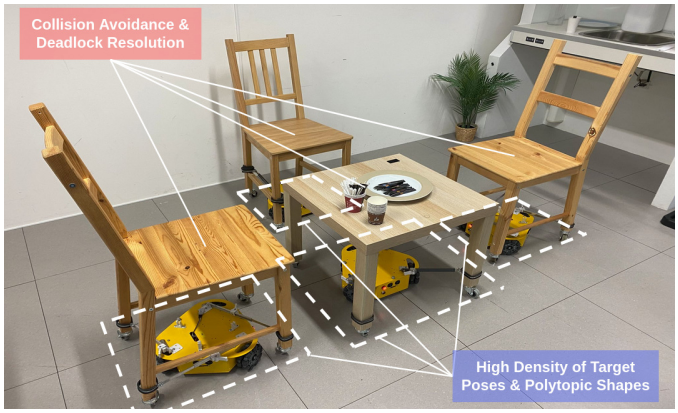

Fig. 1: Motivating application of dense coordination algorithm: mobile assistive furniture swarm.

Challenges occur when coordinating a mobile furniture swarm into dense configurations. As shown in Fig. 1, the formation of furniture in apartment settings is typically confined to limited areas, causing a high density of target poses. Therefore, the convergence of a robot to the target is significantly affected by the shapes of neighboring robots, which tends to cause collisions, oscillations or deadlocks. To overcome these challenges, we propose a solution, named Velocity Potential Field Modulation (VPFM), to coordinate multiple polytopic robots, including both convex and non-convex shapes, and to better manage their reactivity in dense configurations. In addition, we aim for two assistive use cases: (i) the user can define formations of mobile furniture pieces via a tablet interface and let them rearrange autonomously; (ii) the target pose of mobile furniture can dynamically move aside, so that the mobile furniture can clear the way for the wheelchair user.

II. Related Work & Contributions

Multi-robot coordination has been intensely studied for over three decades with various categories of approaches proposed, such as sampling-based methods, velocity obstacle, optimization-based methods, learning-based methods, potential fields and dynamical system modulation.

1) Sampling-based Methods: Sampling-based algorithms construct feasible paths by randomly sampling points from the obstacle-free space and connecting these points into a

Manuscript received 6 April 2025; accepted 15 May 2025. This paper was recommended for publication by Editor M. Ani Hsieh upon evaluation of the reviewers' comments. This work is supported by the project IMEC – EPFL. (*Corresponding author: Lixuan Tang.*)

Lixuan Tang, David Rüegg and Auke Ijspeert are with Biorobotics Laboratory, Swiss Federal Institute of Technology in Lausanne - EPFL, Switzerland (email: lixuan.tang@epfl.ch; david.ruegg.0504@gmail.com; auke.ijspeert@epfl.ch)

Runze Zhang is with Department of Computer Science and Technology, Tongji University, Shanghai, China (email: zhangrunze0311@gmail.com)

Anastasia Bolotnikova is with Robotics and InteractionS Team, Laboratory for Analysis and Architecture of Systems (LAAS-CNRS), France. (email: anastasia.bolotnikova@laas.fr)

Jan Rabaey is with System-Technology Co-Optimization (STCO), IMEC, Belgium. (email: jan.rabaey@imec.be)

Source code of this work can be accessed at https://github.com/Mikasatlx/VPFM-BioRob-EPFL.

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

graph. The Probabilistic Roadmap (PRM) method [1] checks graph edges for potential collisions and computes the shortest collision-free path, which is adapted to multi-robot setups by reducing redundant collision checks. The Rapidly-Exploring Random Trees (RRT) method [2] incrementally explores the space using a space-filling tree and handles inter-robot constraints through message passing. However, sampling-based methods tend to be computationally expensive in multi-robot systems. Moreover, they typically simplify robots as points or small circles without considering specific shapes.

2) Velocity Obstacle: Velocity Obstacle (VO) [3] with its variants (RVO [4], HRVO [5]) calculates a set of velocities that may lead to collisions and ensures that the robot selects a velocity outside this set for execution. A recent variant [6] constructs VO as a function of vertices and robot states for distributed collision avoidance of polytopic robots, referred to as **Polytopic-VO** in this work. However, its experiments only used convex robots with relatively small sizes.

3) Optimization-based Methods: Solutions in this category formulate trajectory generation and collision avoidance as a unified optimization problem. Model Predictive Control (MPC) is regularly applied to coordinate multiple unmanned aerial vehicles simplified as spheres [7][8]. Voronoi Cells are introduced in [9][10][11], where each small circular robot is assigned to a cell disjoint from one another, plans a next waypoint within this cell closest to the target, and updates its cell in a receding horizon manner, only requiring positional information. Considering collision avoidance of polytopic robots, duality-based methods [12][13] transform non-differentiable collision avoidance constraints between polytopes into smooth and differentiable ones. Dual-form constraints are formulated as Discrete-time Control Barrier Functions (DCBF) constraints in [14], which is integrated into MPC, referred to as **CBF-MPC** in this work. Nevertheless, its experiments focus solely on single-robot navigation with obstacle avoidance. In general, as the number of robots increases, optimization-based approaches face a much heavier computational burden.

4) Learning-based Methods: Recent advances in machine learning facilitate its use in multi-robot coordination, which shifts motion planning computations to an offline training phase. Methods such as [15][16][17] employ Multi-Agent Reinforcement Learning (MARL) to learn policies through rewards. This paradigm is currently particularly popular as it does not require large training datasets. However, most learning-based methods simplify robots using basic geometries (e.g., circles or rectangles) and assume relatively small sizes, meaning that the impact of specific shapes is rarely examined.

5) Potential Fields: Artificial Potential Field (APF) [18] guides a point mobile robot or the end-effector of a manipulator to avoid collisions with obstacles in complex shapes by combining attractive forces towards goals and repulsive forces away from obstacles, where the goals are formulated as **attractors** and obstacles as **repulsors**. Velocity Potential Field (**VPF**) [19] models attractions and repulsions as velocities instead of forces to avoid considering dynamic issues. APF methods have been generalized to coordinate multiple robots with convex or non-convex shapes [20][21]. However, since repulsion is applied along the normal direction of the obstacle surface, APF suffers from local minima when repulsion confronts attraction, which tends to cause deadlocks. Inspired by the dynamics of incompressible and irrotational fluids around impenetrable obstacles, Harmonic Potential Functions (HPF) [22] can eliminate local minima in APF, where the potential flow towards obstacles is deflected to the tangent direction to flow around the body. This tangent component helps to escape from local minima and allows the target to be the only attractor of the robot in the whole system. However, it requires the dynamics of robots to follow the properties of harmonic functions, resulting in limited feasible robot motions.

6) Dynamical System Modulation: Dynamical System Modulation (**DSM**) [23] reuses ideas from HPF to deflect flows along the normal direction to the tangent direction of obstacles. It provides a closed-form solution for point robots to reach attractors while avoiding collisions with convex or star-shaped concave obstacles without following the properties of harmonic functions. DSM has been applied to coordinate multiple rectangle robots in [24][25], where the dynamical system is evaluated at a set of control points positioned on each robot. Nevertheless, similar to HPF, DSM is less effective for non-point robots, since the modulation to control points does not necessarily respect rigid body kinematics constraint. In addition, when a robot reaches its target, the attractive velocity becomes zero, meaning that it becomes static and nonreactive to other moving robots. It can cause deadlocks in dense configurations where targets are limited in a small region. A safety module described as repulsive VPF was added over DSM in [25] to coordinate multiple rectangle robots, referred to as **Safe-HDSM** in this work. Although it decreases collision rates significantly, the regulation mechanism of VPF still can not make converged robots reactive again.

7) Contributions: To address the limitations of the previously described methods, we propose our **Velocity Potential Field Modulation (VPFM)**. The main contributions are:

- Our proposed method tackles the dense coordination problem of polytopic swarms by leveraging the repulsion along the normal direction for collision avoidance, and utilizing the tangent flow to resolve deadlocks.
- We propose a more effective mechanism to regulate the strength and range of repulsion compared to Safe-HDSM, enabling better management of robot reactivity to further prevent deadlocks and oscillations.
- We adapt VPFM to assistive use cases: (i) room reconfigurations and (ii) facilitating the indoor movement of wheelchair users. We show that VPFM outperforms the state-of-the-art in simulations, and demonstrate its effectiveness using real-world mobile furniture hardware.

III. Velocity Potential Field Modulation for Reactive Decentralized Coordination

We define the dense coordination problem of a polytopic swarm in Sec. III-A. Then we present how the attraction towards attractors is generated in Sec. III-B, how the attractive velocity is modulated to flow around the robot surface in Sec. III-C, and how the repulsion is generated and regulated to prevent collisions while managing reactivity of robots in

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

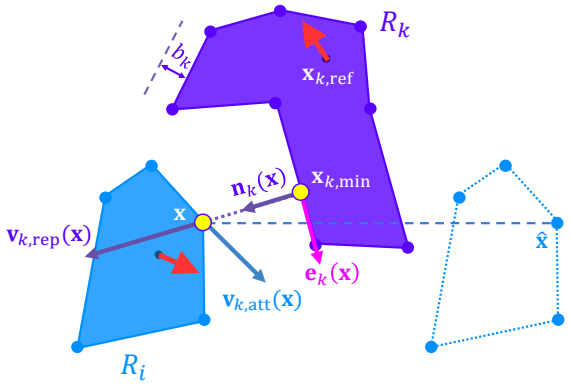

Fig. 2: The attractors of R_i (blue dots on the right) generate attraction to R_i, which is modulated by R_k. R_k itself is also a repulsor which generates repulsion to R_i.

Sec. III-E. A distance function representing the proximity between a point and a robot is introduced in Sec. III-D. Finally, we present how the final linear and angular velocity of a polytopic rigid body is obtained in Sec. III-F. Fig. 2 illustrates the example of the VPFM algorithm principle.

A. Problem Statement

Assume a mobile robot swarm with $N \in \mathbb{N}_+$ robots. We formulate our system in 2D, and assume that all the robots are omni-directional. Each robot $R_i \in \{R_1, R_2, ..., R_N\}$ has a polytopic shape, containing $M_i \in \mathbb{N}_+$ vertices: $\mathbf{X}_i = \{\mathbf{x}_{i,j} \in \mathbb{R}^2 | j = 1, 2, ..., M_i\}$. The shape of R_i can be convex or nonconvex, with the shape area as $S_i \in \mathbb{R}_{>0}$. The target positions (**attractors**) of the vertices set $\mathbf{X}_i$ is denoted as $\hat{\mathbf{X}}_i = \{\hat{\mathbf{x}}_{i,j} \in \mathbb{R}^2 | j = 1, 2, ..., M_i\}$. Using an onboard sensor with sensing radius $r_{i,\text{sen}} \in \mathbb{R}_{>0}$, R_i can sense external states of each surrounding robot R_k, including the shape boundary enclosed by the vertices set $\mathbf{X}_k$, the current position $\mathbf{x}_{k,\text{ref}} \in \mathbb{R}^2$, linear velocity $\dot{\mathbf{x}}_{k,\text{ref}} \in \mathbb{R}^2$ of its reference point (i.e., the centroid of the vertices set $\mathbf{X}_k$), the current orientation $\phi_k \in \mathbb{R}$ and angular velocity $\dot{\phi}_k \in \mathbb{R}$. The internal states of R_k, such as the target position $\hat{\mathbf{x}}_{k,\text{ref}} \in \mathbb{R}^2$ (i.e., the centroid of the attractors set $\hat{\mathbf{X}}_k$), and the target orientation $\hat{\phi}_k \in \mathbb{R}$, are not perceptible by R_i, but are assumed to be shared by local communication between this pair. The user is described as a circle-shaped dynamic nonreactive obstacle whose external states can be sensed by other robots. Assuming that the target poses are distributed in a region with the area as $S_{tot} \in \mathbb{R}_{>0}$, we focus on the dense coordination problem, which means that the density of target poses, defined as $\rho = \frac{\sum_{i=1}^{N} S_i}{S_{tot}}$, is a relatively large value compared to previous work.

B. Attractive Velocity Generation

The kinematics of a polytopic robot is governed by multiple control points. Unlike previous approaches [24][25], control points are now positioned on the boundary of a robot, rather than within its interior. Control points are placed on all vertices of a polygon, and can also be on edges to enable a more precise shape description. As shown in Fig. 2, let $\mathbf{x} \in \mathbb{R}^\eta$ be the position of one control point of R_i, with $\hat{\mathbf{x}} \in \mathbb{R}^\eta$ as its attractor. We set $\eta = 2$ in the 2D scenario. An attractive linear dynamical system that leads $\mathbf{x}$ to its attractor is defined as:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) \quad \text{with} \quad \mathbf{f}(\mathbf{x}) = -(\mathbf{x} - \hat{\mathbf{x}}) \tag{1}$$

where the direction of this attractive velocity is shown as the blue dash line from $\mathbf{x}$ to $\hat{\mathbf{x}}$ in Fig. 2.

C. Attractive Velocity Modulation

As shown in Fig. 2, R_k enters the sensing range $r_{i,\text{sen}}$ of R_i. R_k modulates the initial dynamical system (1) of $\mathbf{x}$, which is compressed (or reversed) along the normal direction, and is stretched along the tangent direction, resulting in the modulated attractive velocity $\mathbf{v}_{k,\text{att}}(\mathbf{x})$ of $\mathbf{x}$:

$$\mathbf{v}_{k,\text{att}}(\mathbf{x}) = \mathbf{M}_k(\mathbf{x})(\mathbf{f}(\mathbf{x}) - \dot{\tilde{\mathbf{x}}}_k) + \dot{\tilde{\mathbf{x}}}_k \tag{2}$$

where $\tilde{\mathbf{x}}_k \in \mathbb{R}^2$ is assumed to be a virtual external point of R_k at the same position as $\mathbf{x}$, and $\dot{\tilde{\mathbf{x}}}_k \in \mathbb{R}^2$ is determined by the rigid body kinematics constraint of R_k at $\tilde{\mathbf{x}}_k = \mathbf{x}$:

$$\dot{\tilde{\mathbf{x}}}_k = \dot{\mathbf{x}}_{k,\text{ref}} + \dot{\phi}_k \times (\tilde{\mathbf{x}}_k - \mathbf{x}_{k,\text{ref}}) \tag{3}$$

with $\mathbf{x}_{k,\text{ref}} \in \mathbb{R}^2$, $\dot{\mathbf{x}}_{k,\text{ref}} \in \mathbb{R}^2$ and $\dot{\phi}_k \in \mathbb{R}$ as the reference point, linear velocity and angular velocity of R_k. Therefore, $\mathbf{f}(\mathbf{x}) - \dot{\tilde{\mathbf{x}}}_k$ is the relative velocity of the initial dynamical system at $\mathbf{x}$ with respect to R_k, which is modulated by the modulation matrix $\mathbf{M}_k(\mathbf{x})$ associated with R_k:

$$\mathbf{M}_k(\mathbf{x}) = \mathbf{E}_k(\mathbf{x})\mathbf{D}_k(\mathbf{x})\mathbf{E}_k(\mathbf{x})^{-1} \tag{4}$$

where $\mathbf{E}_k(\mathbf{x})$ is defined as an orthonormal basis matrix:

$$\mathbf{E}_k(\mathbf{x}) = [\mathbf{n}_k(\mathbf{x}) \quad \mathbf{e}_k(\mathbf{x})] \tag{5}$$

$\mathbf{n}_k(\mathbf{x}) \in \mathbb{R}^2$ is the normal vector from $\mathbf{x}_{k,\text{min}}$ to $\mathbf{x}$, where $\mathbf{x}_{k,\text{min}} \in \mathbb{R}^2$ on R_k has the shortest distance to the control point $\mathbf{x}$. $\mathbf{e}_k(\mathbf{x}) \in \mathbb{R}^2$ is the tangent vector perpendicular to $\mathbf{n}_k(\mathbf{x})$. The diagonal eigenvalue matrix $\mathbf{D}_k(\mathbf{x})$ is defined as:

$$\begin{aligned} &\mathbf{D}_k(\mathbf{x}) = \mathbf{diag}(\lambda_{\mathbf{n}_k}(\mathbf{x}), \lambda_{\mathbf{e}_k}(\mathbf{x})) \quad \text{with} \\ &\lambda_{\mathbf{n}_k}(\mathbf{x}) = 1 - \frac{c_\mathbf{n}}{\Gamma_k(\mathbf{x})}, \lambda_{\mathbf{e}_k}(\mathbf{x}) = 1 + \frac{c_\mathbf{e}}{\Gamma_k(\mathbf{x})} \end{aligned} \tag{6}$$

where $\Gamma_k(\mathbf{x})$ is the distance function associated with R_k that represents the proximity between its surface and the control point $\mathbf{x}$. Γ_k can be defined in various ways according to application requirements, but should satisfy $\Gamma_k(\mathbf{x}) = 1$ for control points reaching the surface of R_k, and increases monotonically as the shortest distance $d_k = ||\mathbf{x} - \mathbf{x}_{k,\text{min}}||$ to R_k increases. $c_\mathbf{n} \in \mathbb{R}_{>0}$ and $c_\mathbf{e} \in \mathbb{R}_{>0}$ control the modulation strength along the normal and tangent directions, respectively. $\lambda_{\mathbf{n}_k}(\mathbf{x}) < 1$ implies the compression strength along the normal direction $\mathbf{n}_k(\mathbf{x})$. $\lambda_{\mathbf{e}_k}(\mathbf{x}) > 1$ implies the strength of stretching along the tangent $\mathbf{e}_k(\mathbf{x})$. As a result, the modulated attractive velocity $\mathbf{v}_{k,\text{att}}(\mathbf{x})$ deviates from its original direction to flow around R_k, as shown in Fig. 2.

Considering there are $N_\text{local} \in \mathbb{N}_+$ robots stepping into the sensing range $r_{i,\text{sen}}$ of R_i, the attractive velocity in (1) is modulated for each robot individually using (2). The final modulated attractive velocity is obtained by a weighted sum:

$$\mathbf{v}_{\text{avg,att}}(\mathbf{x}) = \frac{\sum_{k=1}^{N_\text{local}} \tilde{w}_k \mathbf{v}_{k,\text{att}}(\mathbf{x})}{\sum_{k=1}^{N_\text{local}} \tilde{w}_k} \tag{7}$$

where $\tilde{w}_k$ is the dynamic weight with respect to R_k measured at the control point $\mathbf{x}$ of R_i using Γ_k function:

$$\tilde{w}_k = \frac{r_{i,\text{sen}}}{\Gamma_k(\mathbf{x}) - 1} - 1, k = 1, 2, ..., N_\text{local} \tag{8}$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

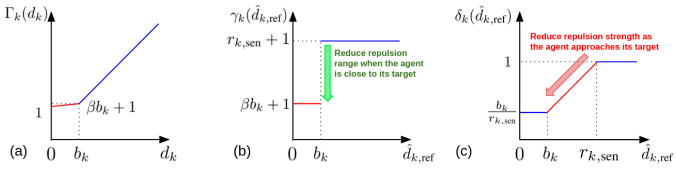

Fig. 3: Profiles of (a) distance function with buffer region Γ_k, (b) local repulsion range γ_k, and (c) repulsion strength regulator δ_k.

D. Distance Function with Buffer Region

Different from the distance function defined in [24][25], we propose a new distance function $\Gamma_k(\mathbf{x}) = \Gamma_k(d_k(\mathbf{x})) = \Gamma_k(d_k)$ to adapt to the dense coordination problem of a polytopic swarm. It introduces a buffer region around the surface of R_k (purple dash line in Fig. 2) with size $b_k \in \mathbb{R}_{>0}$:

$$\Gamma_k(d_k) = \begin{cases} \beta d_k + 1, & 0 \le d_k < b_k \\ d_k + (\beta - 1)b_k + 1, & d_k \ge b_k \end{cases} \tag{9}$$

$\beta \in \mathbb{R}_{>0} \to 0^+$ is a coefficient with small value. The profile of $\Gamma_k(d_k)$ is shown in Fig. 3(a). When $0 \le d_k < b_k$, the control point $\mathbf{x}$ falls into the buffer region with $\Gamma_k(d_k) \to 1^+$, which is shown as the red segment in Fig. 3(a). This suggests a high risk of collision within the buffer region, where the attractive velocity modulation will consistently remain strong, according to (6). When $d_k \ge b_k$, the control point $\mathbf{x}$ is outside the buffer region, and the increase of $\Gamma_k(d_k)$ is proportional to the increase of d_k, which is shown as the blue segment in Fig. 3(a). In this case, $\Gamma_k(d_k)$ can directly reveal the real distance value d_k. Note that, when $d_k = 0$, it will be justified as a failure case due to collision, and we assume that all robots within the system will terminate immediately in this work.

E. Repulsive Velocity Generation and Regulation

The repulsive velocity potential is introduced to avoid collision along the normal direction, and to keep robots reactive even if they reach targets. As shown in Fig. 2, we treat R_k as a repulsor that sends repulsion to the control point $\mathbf{x}$ of R_i along the normal direction $\mathbf{n}_k(\mathbf{x})$. However, excessive strength or range of repulsion can lead to oscillations and hinder robots from converging. An attractor-repulsor coupling mechanism is proposed, in which both the strength and the range of repulsion exerted by R_k decrease as it gradually approaches its attractor $\hat{\mathbf{x}}_{k,\text{ref}}$, mitigating the risk of oscillations and deadlocks. The repulsive velocity $\mathbf{v}_{k,\text{rep}}(\mathbf{x})$ from R_k to $\mathbf{x}$ is defined as:

$$\mathbf{v}_{k,\text{rep}}(\mathbf{x}) = \frac{\delta_k}{r_{k,\text{sen}}(\Gamma_k(d_k) - 1)}\mathbf{n}_k(\mathbf{x}) \tag{10}$$

(10) is applied only if the control point $\mathbf{x}$ of R_i falls into the local repulsion range of R_k: $\Gamma_k(d_k) < \gamma_k$. The local repulsion range γ_k of R_k is regulated by the direct distance from R_k to its attractor $\hat{d}_{k,\text{ref}} = ||\mathbf{x}_{k,\text{ref}} - \hat{\mathbf{x}}_{k,\text{ref}}||$:

$$\gamma_k(\hat{d}_{k,\text{ref}}) = \begin{cases} r_{k,\text{sen}} + 1, & \hat{d}_{k,\text{ref}} > b_k \\ \beta b_k + 1, & 0 \le \hat{d}_{k,\text{ref}} \le b_k \end{cases} \tag{11}$$

The profile of $\gamma_k(\hat{d}_{k,\text{ref}})$ is shown as Fig. 3(b). When $\hat{d}_{k,\text{ref}} > b_k$, the local repulsion range reveals the local sensing range $r_{k,\text{sen}}$ (the blue segment in Fig. 3(b)). When $0 \le \hat{d}_{k,\text{ref}} \le b_k$, R_k is very close to its attractor, and $\gamma_k(\hat{d}_{k,\text{ref}}) = \beta b_k + 1$ is the Γ_k value achieved on the boundary of the buffer region $d_k = b_k$, according to (9). This means that the local repulsion range of R_k reduces to its buffer region when it nearly converges. In this case, repulsion will be applied only if the control point steps inside the buffer region. δ_k in (10) is used to regulate the repulsion strength of R_k, also relying on $\hat{d}_{k,\text{ref}}$:

$$\delta_k(\hat{d}_{k,\text{ref}}) = \begin{cases} 1, & \hat{d}_{k,\text{ref}} > r_{k,\text{sen}} \\ \frac{\hat{d}_{k,\text{ref}}}{r_{k,\text{sen}}}, & b_k \le \hat{d}_{k,\text{ref}} \le r_{k,\text{sen}} \\ \frac{b_k}{r_{k,\text{sen}}}, & 0 \le \hat{d}_{k,\text{ref}} < b_k \end{cases} \tag{12}$$

The profile of $\delta_k(\hat{d}_{k,\text{ref}})$ is shown as Fig. 3(c). Within range $\hat{d}_{k,\text{ref}} \in [b_k, r_{k,\text{sen}}]$, δ_k value shrinks linearly as R_k is gradually approaching its target, which is shown as the red segment in Fig. 3(c). The combined effect of (11) and (12) ensures that repulsions from R_k will be significant if it is still far from its attractor $\hat{\mathbf{x}}_{k,\text{ref}}$. As R_k approaches $\hat{\mathbf{x}}_{k,\text{ref}}$, both the strength and the range of repulsions shrink, so that other robots that have not reached their targets are less likely to be blocked. Both the discrete switching of the repulsion range and the continuous varying of the repulsion strength prevent velocities from jittering. Since $\hat{\mathbf{x}}_{k,\text{ref}}$ and $r_{k,\text{sen}}$ are not perceptible by R_i, local communication is required to exchange information in order to compute $\delta_k(\hat{d}_{k,\text{ref}})$ for decentralized coordination.

The control point $\mathbf{x}$ of R_i might receive repulsions from multiple robots, but only the repulsive velocity with the maximal value $\mathbf{v}_{\max,\text{rep}}(\mathbf{x})$ is considered. The velocity $\dot{\mathbf{x}}$ evaluated for $\mathbf{x}$ is defined as the sum of the modulated attractive velocity $\mathbf{v}_{\text{avg,att}}(\mathbf{x})$ and the repulsive velocity $\mathbf{v}_{\max,\text{rep}}(\mathbf{x})$:

$$\dot{\mathbf{x}} = \mathbf{v}_{\text{avg,att}}(\mathbf{x}) + \mathbf{v}_{\max,\text{rep}}(\mathbf{x}) \tag{13}$$

F. Rigid Body Constraint via Weighted Least Squares

In previous sections, velocities of different control points are evaluated separately, which do not necessarily respect rigid body constraints (i.e., they would lead to deformations of the body). Therefore, a weighting scheme of control point velocities is required to compute the final linear velocity $\dot{\mathbf{x}}_{i,\text{ref}}$ and the angular velocity $\dot{\phi}_i$ of R_i. Assume that the jth control point of R_i is $\mathbf{x}_{i,j} \in \mathbb{R}^2, j = 1, 2...M_i$, and that there are $N_{\text{local}} \in \mathbb{N}_+$ robots entering the sensing range $r_{i,\text{sen}}$ of R_i. To ensure that the control point closer to surrounding robots has a stronger influence, we define $w_{i,j}$ as the weight of $\mathbf{x}_{i,j}$:

$$w_{i,j} = \max\{\tilde{w}_{i,j,k} | k = 1, 2, ...N_{\text{local}}\} \tag{14}$$

where the definition of $\tilde{w}_{i,j,k}$ is given in (8). The final control point velocity $\dot{\mathbf{x}}_{i,j}$ should follow the rigid body constraint:

$$\dot{\mathbf{x}}_{i,j} = \dot{\mathbf{x}}_{i,\text{ref}} + \dot{\phi}_i \times (\mathbf{x}_{i,j} - \mathbf{x}_{i,\text{ref}}) \tag{15}$$

which can be rewritten in the form of $\mathbf{A}_{i,j}\xi_i = \mathbf{b}_{i,j}$:

$$\underbrace{\begin{bmatrix} 1 & 0 & -(y_{i,j} - y_{i,\text{ref}}) \\ 0 & 1 & x_{i,j} - x_{i,\text{ref}} \end{bmatrix}}_{\mathbf{A}_{i,j} \in \mathbb{R}^{2\times 3}} \underbrace{\begin{bmatrix} \dot{x}_{i,\text{ref}} \\ \dot{y}_{i,\text{ref}} \\ \dot{\phi}_i \end{bmatrix}}_{\xi_i \in \mathbb{R}^{3\times 1}} = \underbrace{\begin{bmatrix} \dot{x}_{i,j} \\ \dot{y}_{i,j} \end{bmatrix}}_{\mathbf{b}_{i,j} \in \mathbb{R}^{2\times 1}} \tag{16}$$

When stacking (16) for M_i control points of R_i, $2M_i$ equations in total are provided with the number of unknown variables being 3, namely $\dot{x}_{i,\text{ref}}$, $\dot{y}_{i,\text{ref}}$ and $\dot{\phi}_i$, which can

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

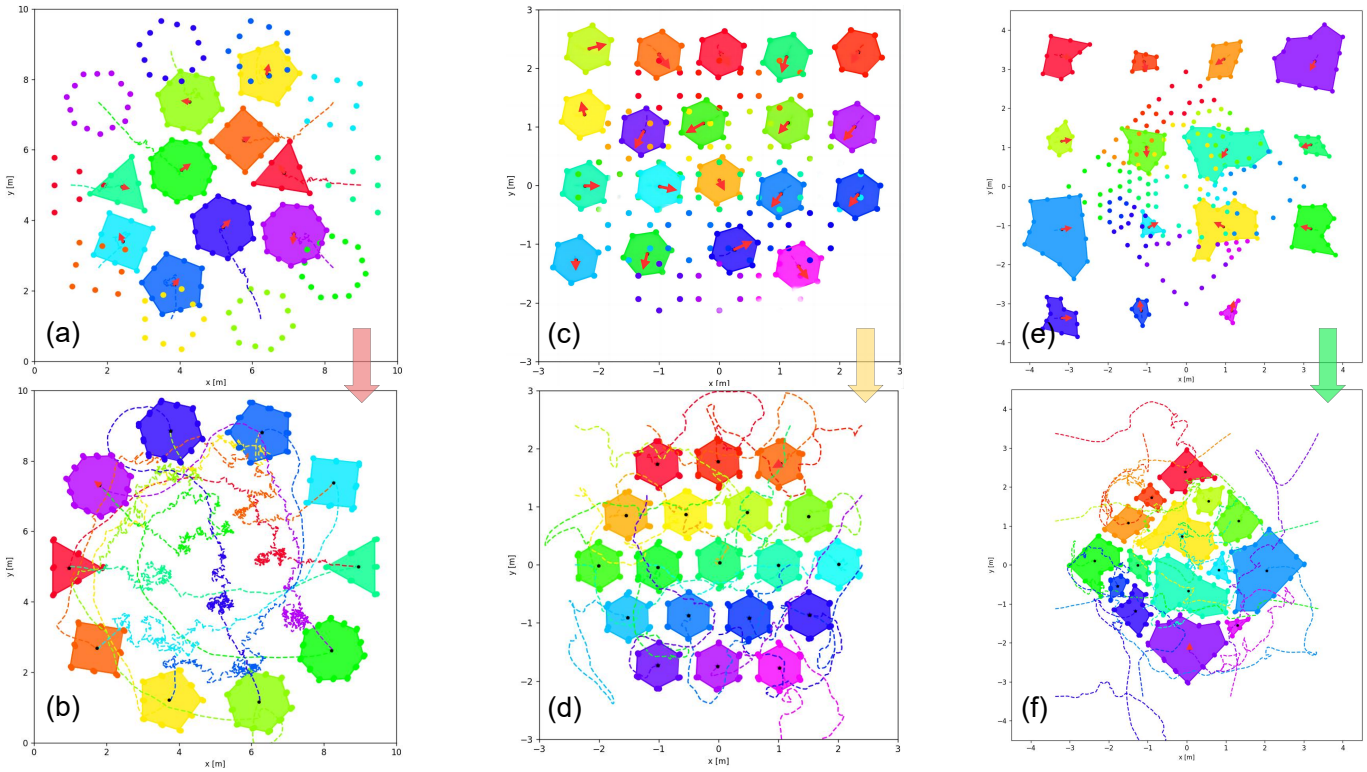

Fig. 4: Qualitative results of Antipodal Position Switching of regular polygons (a)(b), Dense Reconfiguration of regular hexagons (c)(d), and Auto-assembly of irregular polytopic parts (e)(f). Each robot has to reach its target with the same color (also applies to Fig. 5).

Radius of Regular Polygons / Density of Target Poses	Method	Scenario Convergence Rate (%)	Collision Rate (%)	Deadlock Rate (%)	Convergence Time (s)	Distance Ratio (%)	Time Per Iteration (s)
0.1 m / 0.22%	**Polytopic-VO**	**100**	**0**	**0**	17.1 ± 0.6	**102 ± 1**	**0.014 ± 0.001**
	CBF-MPC	84	16	**0**	**7.4** ± 0.6	111 ± 3	2.025 ± 0.102
	DSM	99	**0**	1	21.5 ± 2.7	122 ± 8	0.144 ± 0.028
	VPF	96	**0**	4	19.9 ± 1.5	123 ± 6	0.087 ± 0.002
	Safe-HDSM	**100**	**0**	**0**	25.7 ± 2.2	125 ± 5	0.141 ± 0.008
	VPFM	**100**	**0**	**0**	20.2 ± 1.5	123 ± 4	0.141 ± 0.007
0.9 m / 17.83%	Polytopic-VO	9	90	1	**45.7 ± 13.8**	**150 ± 10**	**0.039 ± 0.008**
	CBF-MPC	0	22	78	-	-	2.239 ± 0.253
	DSM	9	88	3	77.9 ± 10.2	343 ± 37	0.148 ± 0.003
	VPF	95	5	**0**	53.0 ± 11.4	311 ± 61	0.088 ± 0.008
	Safe-HDSM	74	**0**	26	46.5 ± 10.5	207 ± 18	0.149 ± 0.006
	VPFM	**99**	1	**0**	50.0 ± 7.5	299 ± 39	0.143 ± 0.008

TABLE I: Quantitative evaluation of Antipodal Position Switching of regular polygons with varying radius

be formulated as a Least Square problem $\mathbf{A}_i\xi_i = \mathbf{b}_i$ with $\mathbf{A}_i \in \mathbb{R}^{2M_i\times 3}$, $\xi_i \in \mathbb{R}^{3\times 1}$ and $\mathbf{b}_i \in \mathbb{R}^{2M_i\times 1}$.

It is further transformed into a Weighted Least Square problem by introducing the weight matrix $\mathbf{W}_i \in \mathbb{R}^{2M_i\times 2M_i}$:

$$(\mathbf{W}_i\mathbf{A}_i)\,\xi_i = \mathbf{W}_i\mathbf{b}_i \tag{17}$$

The weight matrix is a diagonal matrix whose element $w_{i,j}$ is computed by (14):

$$\mathbf{W} = \mathbf{diag}\left([w_{i,1}\ w_{i,1}\ \ldots\ w_{i,j}\ w_{i,j}\ \ldots\ w_{i,M_i}\ w_{i,M_i}]\right) \tag{18}$$

The results of the Weighted Least Squares serve as the final linear velocity $\dot{\mathbf{x}}_{i,\text{ref}}$ and the angular velocity $\dot{\phi}_i$ of R_i.

IV. Simulation Experiments

In this section, we first introduce evaluation metrics used in simulations, then show results from multiple experimental setups. All simulations were executed on a PC with Intel(R) Core(TM) i7-13700 CPU using Python 3.8.

A. Evaluation Metrics

The convergence of a robot is defined as the reaching of its target within tolerable error ($0.05\,\text{m}$ for position and $0.05\,\text{rad}$ for orientation) and within a given time budget $T_{\max}$. Each simulation setup is repeated for 100 times with randomized initial conditions (see Sec. IV-B). We report the following metrics: **(i) Collision Rate:** If a robot collides with another object, this run will be terminated immediately, and counted as a collision. We report the ratio of collisions among 100 runs. **(ii) Deadlock Rate:** If there still exists any robot that has not yet converged to its target at $T_{\max}$, this run will be counted as a deadlock. We report the ratio of deadlocks among 100 runs. **(iii) Scenario Convergence Rate:** If all robots converge within $T_{\max}$, this run will be counted as a scenario convergence. We report the ratio of scenario convergences among 100 runs. **(iv) Robot Convergence Rate:** At the end of each run, we calculate the ratio of converged robots among all robots. We report the average ratio for 100 runs.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Radius of Hexagons / Density of Target Poses	Method	Robot Convergence Rate (%)	Collision Rate (%)	Deadlock Rate (%)	Convergence Time (s)	Distance Ratio (%)	Time Per Iteration (s)
0.1 m / 1.37%	Polytopic-VO	86.7	**0**	89	47.8 ± 12.3	869 ± 266	**0.038 ± 0.007**
	CBF-MPC	97.9	8	12	15.8 ± 2.2	**114 ± 3**	9.661 ± 0.361
	DSM	**100**	**0**	**0**	10.6 ± 1.8	119 ± 6	0.302 ± 0.016
	VPF	**100**	**0**	**0**	**9.9 ± 1.0**	160 ± 23	0.189 ± 0.005
	Safe-HDSM	97.3	**0**	39	17.6 ± 4.4	183 ± 23	0.313 ± 0.008
	VPFM	**100**	**0**	**0**	11.2 ± 1.3	174 ± 25	0.297 ± 0.006
0.4 m / 21.94%	Polytopic-VO	61.1	**0**	100	-	-	**0.063 ± 0.012**
	CBF-MPC	11.8	63	37	-	-	10.063 ± 0.357
	DSM	10.8	87	13	-	-	0.331 ± 0.007
	VPF	86.2	1	15	25.9 ± 7.3	**376 ± 80**	0.191 ± 0.002
	Safe-HDSM	40.2	**0**	100	-	-	0.340 ± 0.004
	VPFM	**96.6**	1	**4**	**22.0 ± 7.8**	397 ± 74	0.315 ± 0.004

TABLE II: Quantitative evaluation of Dense Reconfiguration of hexagons with varying radius

Method	Robot Convergence Rate (%)	Collision Rate (%)	Deadlock Rate (%)	Convergence Time (s)	Distance Ratio (%)	Time Per Iteration (s)
DSM	9.4	96	**4**	-	-	0.333 ± 0.014
VPF	55.2	3	57	44.2 ± 21.8	391 ± 152	**0.199 ± 0.009**
Safe-HDSM	45.5	1	99	-	-	0.345 ± 0.012
VPFM	**96.1**	**0**	6	**28.0 ± 13.6**	**369 ± 149**	0.332 ± 0.007

TABLE III: Quantitative evaluation of Auto-assembly of irregular polytopic parts

(v) Convergence Time / Std: If one run reaches a scenario convergence, we count the time spent achieving convergence (number of iterations times the timestep Δt). We report the average convergence time with its standard deviation among all scenario converged runs. **(vi) Distance Ratio / Std:** If one run reaches a scenario convergence, we record the average ratio of travel distance divided by the straight line distance between the initial and target position over all robots. We report the average of this value with its standard deviation among all scenario converged runs. **(vii) Time Per Iteration / Std:** We record the average execution time per iteration for each run (not included in computing Convergence Time), and report the average with its standard deviation among 100 runs.

B. Simulation Results

1) Evaluation of Antipodal Position Switching: The setup is visualized in Fig. 4(a)(b), where the target poses are distributed in a $10\,\mathrm{m} \times 10\,\mathrm{m}$ area. All robots are initialized evenly on a circle, with their targets set on the opposite side. They meet at the center and try to avoid collision with each other. The radius of the formation circle is $4\,\mathrm{m}$. Each robot is a regular polygon with number of vertices varying from 3 to 7. We place one control point at each vertex and at each edge. The number of polygons is fixed to 10, and we increase the radius of regular polygons from $0.1\,\mathrm{m}$ to $0.9\,\mathrm{m}$, with the size of buffer region $b = 0.3\,\mathrm{m}$. The maximal linear velocity of each robot is randomized within range $[0.5\,\mathrm{m\,s^{-1}},\ 1.0\,\mathrm{m\,s^{-1}}]$. We set $T_{\max} = 120\,\mathrm{s}$ and $\Delta t = 0.1\,\mathrm{s}$. For our VPFM algorithm, we set $r_{\text{sen}} = 2\,\mathrm{m}$, and $c_{\mathbf{n}} = c_{\mathbf{e}} = 1$.

The quantitative results are shown in Table I. When the target pose density is 0.22%, all methods except for CBF-MPC can achieve near 100% in Scenario Convergence Rate, within which Polytopic-VO is the most efficient for all metrics. As the target pose density increases to 17.83%, only VPF and VPFM can maintain a high Scenario Convergence Rate. Polytopic-VO and DSM suffer from collisions, whereas CBF-MPC and Safe-HDSM suffer from deadlocks. CBF-MPC has the highest Time Per Iteration. The result shows that both VPF and VPFM are effective in pure collision avoidance for polytopic robots with relatively large sizes.

2) Evaluation of Dense Reconfiguration: The setup with its qualitative result is shown in Fig. 4(c)(d), where the target poses are distributed in a $6\,\mathrm{m} \times 6\,\mathrm{m}$ area. The number of regular hexagons is fixed to 19, and we increase the radius from $0.1\,\mathrm{m}$ to $0.4\,\mathrm{m}$, with the buffer region size $b = 0.15\,\mathrm{m}$. We place one control point at each vertex. The target pose of each hexagon serves as a unit within a compact formation. To randomly initialize the poses of hexagons, the workspace is divided into a 5×5 grid. Each hexagon is placed into each grid cell in random order. The maximal linear velocity of each robot is set as $0.6\,\mathrm{m\,s^{-1}}$. We set $\Delta t = 0.05\,\mathrm{s}$, with $T_{\max}$, r_{sen}, $c_{\mathbf{n}}$ and $c_{\mathbf{e}}$ the same as in Sec. IV-B1.

The quantitative result is shown as Table II. When the density of target poses is 1.37%, all methods except for Polytopic-VO can achieve near 100% in Robot Convergence Rate, with DSM, VPF and VPFM on par with each other. Although Polytopic-VO is the fastest in Time Per Iteration, it has high Deadlock Rates, Convergence Time and Distance Ratio. This means that robots continue to influence each other using the VO protocol even when they are close to converging on targets, making them oscillate back and forth. Safe-HDSM also encounters this issue, but its high deadlock rate is primarily due to robots becoming nonreactive upon reaching their targets while still exerting repulsive forces on other robots in motion. When the target pose density increases to 21.94%, Polytopic-VO and Safe-HDSM still suffer from severe deadlocks, and CBF-MPC and DSM exhibit high Collision Rates. As an optimization-based method, CBF-MPC is the slowest in Time Per Iteration. Compared to VPF, the additional velocity modulation in VPFM helps to resolve deadlocks by providing tangent velocity flowing around the surface, demonstrating that our method is an effective solution to the dense coordination problem of polytopic swarms when a compact formation is required.

3) Evaluation of Auto-assembly: The setup with its qualitative result is visualized in Fig. 4(e)(f). 15 irregular polytopic parts are commanded to assemble in a compact square-like structure, with buffer region size $b = 0.18\,\mathrm{m}$. The random initialization of poses, T_{max}, Δt, r_{sen}, $c_{\mathbf{n}}$ and $c_{\mathbf{e}}$ are the same as in Sec. IV-B2. The quantitative result is shown as Table III. Collision Rate reaches 96% for DSM. Both VPF and Safe-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

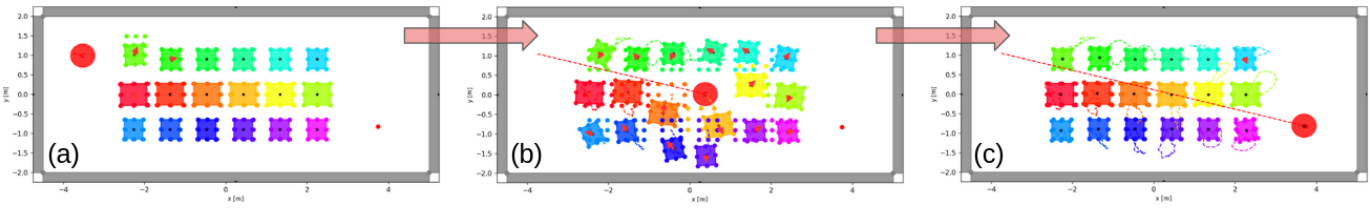

Fig. 5: Qualitative results of dynamic attractors for clearing the way for the user.

Attractor Status	Static	Dynamic
Scenario Convergence Rate (%)	72	**91**
Collision Rate (%)	28	**9**
Deadlock Rate (%)	**0**	**0**

TABLE IV: Quantitative results of adapting to dynamic attractors

Use Case Setup	Scenario Convergence Rate (%)	Time Per Iteration (s)
Room Reconfiguration	100 ± 0.0	0.073 ± 0.001
Room Reconfiguration with Disturbance	100 ± 0.0	0.074 ± 0.001
Clearing the Way for the User	100 ± 0.0	0.073 ± 0.001

TABLE V: Quantitative evaluation of real-world experiments

HDSM suffer from deadlocks. Compared to VPF, additional velocity modulation helps VPFM to resolve deadlocks, but causes more computational burden. Compared to Safe-HDSM, the attractor-repulsor coupling mechanism in VPFM regulates the repulsive potential more properly, ensuring that robots maintain reactivity even after reaching their targets. The result demonstrates that VPFM can coordinate a polytopic robot swarm with irregular shapes (both convex and non-convex) and different sizes. Nevertheless, the density of control points should align with that of the smallest robot, which tends to cause heavier computation when coordinating a swarm of robots with significant size disparities.

4) Evaluation of Clearing the Way for the User: In this section, we evaluate the performance of VPFM when adapting to dynamic attractors. For this experiment, the target pose of a robot moves perpendicularly to the velocity of a user approaching. When following the motion of the attractor, the robot actually clears the path for the user, as shown in Fig. 5(b). The state machine that controls the motion of the attractors is similar to the previous work [24]. The setup with its qualitative result is visualized in Fig. 5. The workspace with size $9.5\,\mathrm{m} \times 4\,\mathrm{m}$ is surrounded by walls. Six mobile tables ($0.6\,\mathrm{m} \times 0.55\,\mathrm{m}$) and 12 mobile chairs ($0.5\,\mathrm{m} \times 0.45\,\mathrm{m}$) are introduced, with the size of buffer region fixed to $0.15\,\mathrm{m}$. The nonreactive user moves along a straight line with constant speed $0.5\,\mathrm{m\,s^{-1}}$, whose initial and target positions are randomized on different sides. We set $T_{\max} = 80\,\mathrm{s}$, with other parameters the same as in Sec. IV-B2. The quantitative result is shown in Table IV. When attractors are static, VPFM can achieve 72% Scenario Convergence Rate, with Collision Rate as 28%. When attractors dynamically move to clear the way, Collision Rate drops to 9%. The result demonstrates that our VPFM adapts well to the dynamic attractor setup.

V. Real-World Experiment

A. Experimental Setup

Fig. 6 shows our real-world setups, with furniture rendered mobile by using Omnibots proposed in our previous work [26]. The maximal linear velocity of Omnibot is set as $0.3\,\mathrm{m\,s^{-1}}$. The workspace with size $4\,\mathrm{m} \times 3\,\mathrm{m}$ is surrounded by virtual walls. Three chairs with size $0.5\,\mathrm{m} \times 0.42\,\mathrm{m}$ and one table with $0.55\,\mathrm{m} \times 0.55\,\mathrm{m}$ are introduced, with the buffer region size as $0.15\,\mathrm{m}$. Due to current hardware limitations, Omnibot is currently not equipped with an onboard sensor. The centralized localization is performed by the Vicon motion capture system, which tracks markers attached to the mobile furniture. The localization results serve as input to VPFM, which is executed on the same centralized PC as in Sec. IV. The outputs are velocity commands sent to all Omnibots over Bluetooth via the Robot Operating System (ROS). We set $T_{\max} = 120\,\mathrm{s}$ and $\Delta t = 0.1\,\mathrm{s}$, with r_{sen}, $c_{\mathbf{n}}$ and $c_{\mathbf{e}}$ the same as all previous sections. We use the same metrics mentioned in Sec. IV-A.

B. Results with Real Robots

Fig. 6(a)(b) visualizes the use case of Room Reconfiguration, where mobile furniture pieces rearrange themselves autonomously from one setup (either structured or randomly scattered) to another user-defined formation. Fig. 6(c)(d) illustrates Room Reconfiguration with Disturbance, where random disturbances (pulling, pushing or stopping) are manually applied to the rearrangement procedure. Fig. 6(e)(f) illustrates the use case of Clearing the Way for the User. The casual movement of the wheelchair user is adapted by other mobile furniture pieces, whose target poses can move perpendicularly to the user's velocity using the same dynamic attractor mechanism as in Sec. IV-B4, shown as the green rectangle labels in Fig. 6(e). All setups are repeated for 10 times respectively. The quantitative results are given in Table V, which demonstrates that our VPFM algorithm is an effective solution to achieve desired assistive use cases, which can be executed in real time, and is also robust to external disturbances. However, we notice that, when the mobile furniture is approaching its target, the value of its velocity commands also nearly decreases to 0, which increases the difficulty to get over friction before arrival.

VI. Conclusion

We present VPFM as an effective solution to the dense coordination problem of a polytopic swarm, aiming to reduce collisions, oscillations, and deadlocks in scenarios requiring a dense target pose configuration. We have adapted this approach for assistive applications to enhance the lives of individuals with limited mobility and validated it through both simulations and real-world experiments. Nevertheless, although the results suggest minimal collisions and deadlocks under taxing circumstances, the convergence of robots to their targets is currently not guaranteed or provable. We hope to establish

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

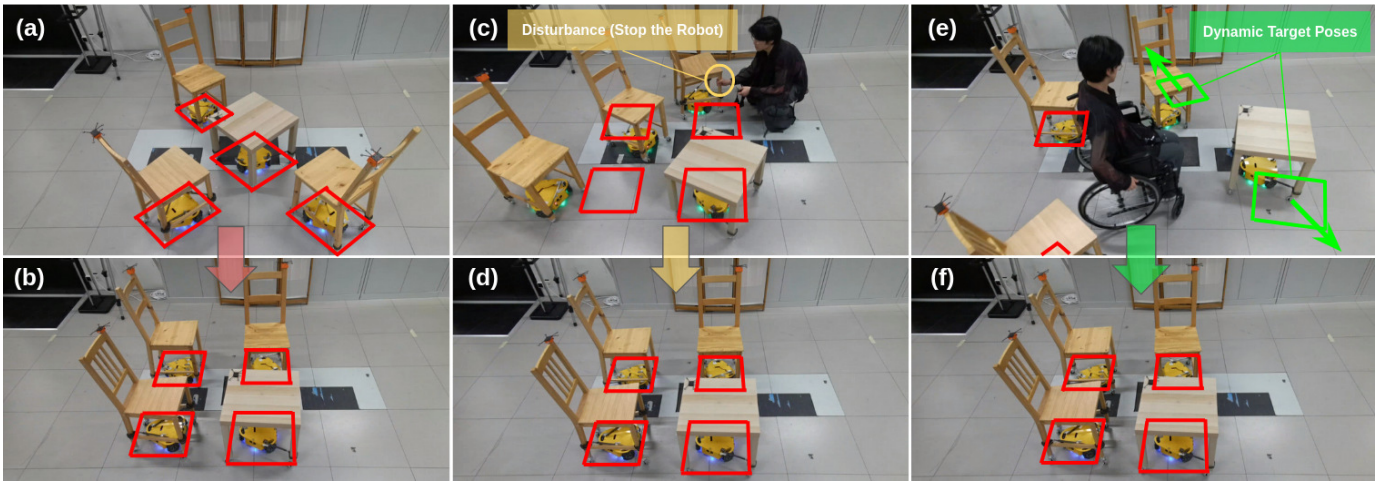

Fig. 6: Qualitative results of real-world experiments. (a)(b) are for Room Reconfiguration. (c)(d) are for Room Reconfiguration with Disturbance. (e)(f) are for Clearing the Way for the User. The target poses of mobile furniture are highlighted using rectangle labels.

bounds in our future research. We also plan to integrate onboard sensing into robot hardware to assess its capability for decentralized coordination, extend the workspace from 2D to 3D, and replace the Vicon system using a markerless solution.

REFERENCES

[1] G. Sánchez and J.-C. Latombe, "A single-query bi-directional probabilistic roadmap planner with lazy collision checking," in *Robotics Research* (R. A. Jarvis and A. Zelinsky, eds.), (Berlin, Heidelberg), pp. 403–417, Springer Berlin Heidelberg, 2003.

[2] V. R. Desaraju and J. P. How, "Decentralized path planning for multi-agent teams in complex environments using rapidly-exploring random trees," in *2011 IEEE International Conference on Robotics and Automation*, pp. 4956–4961, 2011.

[3] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, pp. 760 – 772, 1998.

[4] J. van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, pp. 1928–1935, 2008.

[5] J. Snape, J. v. d. Berg, S. J. Guy, and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 696–706, 2011.

[6] J. Huang, J. Zeng, X. Chi, K. Sreenath, Z. Liu, and H. Su, "Velocity obstacle for polytopic collision avoidance for distributed multi-robot systems," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3502–3509, 2023.

[7] M. Kamel, J. Alonso-Mora, R. Siegwart, and J. Nieto, "Robust collision avoidance for multiple micro aerial vehicles using nonlinear model predictive control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 236–243, 2017.

[8] B. Lindqvist, P. Sopasakis, and G. Nikolakopoulos, "A scalable distributed collision avoidance scheme for multi-agent uav systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 9212–9218, 2021.

[9] D. Zhou, Z. Wang, S. Bandyopadhyay, and M. Schwager, "Fast, on-line collision avoidance for dynamic vehicles using buffered voronoi cells," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1047–1054, 2017.

[10] M. Abdullhak and A. Vardy, "Deadlock prediction and recovery for distributed collision avoidance with buffered voronoi cells," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 429–436, 2021.

[11] S. H. Arul and D. Manocha, "Cglr: Dense multi-agent navigation using voronoi cells and congestion metric-based replanning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7213–7220, 2022.

[12] X. Zhang, A. Liniger, A. Sakai, and F. Borrelli, "Autonomous parking using optimization-based collision avoidance," in *IEEE Conference on Decision and Control (CDC)*, pp. 4327–4332, 2018.

[13] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-based collision avoidance," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 972–983, 2021.

[14] A. Thirugnanam, J. Zeng, and K. Sreenath, "Safety-critical control and planning for obstacle avoidance between polytopes with control barrier functions," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 286–292, 2022.

[15] S. H. Semnani, H. Liu, M. Everett, A. de Ruiter, and J. P. How, "Multi-agent motion planning for dense and dynamic environments via deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3221–3226, 2020.

[16] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10357–10377, 2021.

[17] R. Han, S. Chen, S. Wang, Z. Zhang, R. Gao, Q. Hao, and J. Pan, "Reinforcement learned distributed multi-robot navigation with reciprocal velocity obstacle shaped rewards," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 5896–5903, 2022.

[18] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Proceedings. IEEE International Conference on Robotics and Automation*, vol. 2, pp. 500–505, 1985.

[19] M. Haghshenas Jaryani, "An effective manipulator trajectory planning with obstacles using virtual potential field method," in *IEEE International Conference on Systems, Man and Cybernetics*, pp. 1573–1578, 2007.

[20] C. Warren, "Multiple robot path coordination using artificial potential fields," in *Proceedings., IEEE International Conference on Robotics and Automation*, pp. 500–505 vol.1, 1990.

[21] R. Gayle, W. Moss, M. C. Lin, and D. Manocha, "Multi-robot coordination using generalized social potential fields," in *IEEE International Conference on Robotics and Automation*, pp. 106–113, 2009.

[22] J.-O. Kim and P. Khosla, "Real-time obstacle avoidance using harmonic potential functions," *IEEE Transactions on Robotics and Automation*, vol. 8, no. 3, pp. 338–349, 1992.

[23] L. Huber, A. Billard, and J.-J. Slotine, "Avoidance of convex and concave obstacles with convergence ensured through contraction," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1462–1469, 2019.

[24] F. M. Conzelmann, L. Huber, D. Paez-Granados, A. Bolotnikova, A. Ijspeert, and A. Billard, "A dynamical system approach to decentralized collision-free autonomous coordination of a mobile assistive furniture swarm," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7259–7265, 2022.

[25] L.-N. Douce, A. Menichelli, L. Huber, A. Bolotnikova, D. Paez-Granados, A. Ijspeert, and A. Billard, "Agent prioritization and virtual drag minimization in dynamical system modulation for obstacle avoidance of decentralized swarms," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–7, 2023.

[26] L. Tang, C. Ning, G. Adaimi, A. Ijspeert, A. Alahi, and A. Bolotnikova, "Real-time localization for closed-loop control of assistive furniture," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4799–4806, 2023.