

# Efficient Multi-Agent Global Navigation Using Interpolating Bridges

Liang He<sup>1</sup> and Jia Pan<sup>2</sup> and Dinesh Manocha<sup>1</sup>

**Abstract**—We present a novel approach for collision-free global navigation for continuous-time multi-agent systems with general linear dynamics. Our approach is general and can be used to perform collision-free navigation in 2D and 3D workspaces with narrow passages and crowded regions. As part of pre-computation, we compute multiple bridges in the narrow or tight regions in the workspace using kinodynamic RRT algorithms. Our bridge has certain geometric properties that enable us to calculate a collision-free trajectory for each agent using simple interpolation at runtime. Moreover, we combine interpolated bridge trajectories with local multi-agent navigation algorithms to compute global collision-free paths for each agent. The overall approach combines the performance benefits of coupled multi-agent algorithms with the precomputed trajectories of the bridges to handle challenging scenarios. In practice, our approach can perform global navigation for tens to hundreds of agents on a single CPU core in 2D and 3D workspaces.

## I. INTRODUCTION

Multi-agent navigation algorithms are widely used for motion planning among static and dynamic obstacles. The underlying applications include cooperative surveillance, sensor networks, swarm navigation, and simulation of animated characters or human crowds in games and virtual worlds. One key problem in multi-agent navigation is the computation of collision-free trajectories for agents, given their own initial and goal positions.

At a broad level, prior approaches can be classified into *coupled* or *decoupled* planners. A coupled planner aggregates all the individual robots into one large composite system and leverages classical motion planners (e.g., sampling-based planners) to compute collision-free trajectories for all agents. On the other hand, a decoupled planner computes a trajectory for each robot individually for a short horizon (e.g., a few time-steps), and then performs a velocity coordination to resolve the collision between the local trajectories of all agents. Different techniques have been proposed to compute local collision-free paths or schedule their motion.

Coupled planners are (probabilistically) complete in theory and thus can provide rigorous guarantees about collision avoidance between the agents and the obstacles. However, as the number of agents in the scene increases, the resulting dimension of the system’s configuration space increases linearly and current methods are only practical for a few agents. Decoupled planners are usually faster because fewer degrees of freedom are taken into account at a time. Unfortunately, the velocity coordination may not resolve all collisions, and

the agents may get stuck in crowded scenarios or block each other (e.g., due to the inevitable collision states [1], [2]). Thus, in challenging scenarios with narrow passages, a decoupled planner may take a long time or even be unable to find a solution when one exists. In addition, the velocity coordination can be slow in crowded environments, where all agents have to move toward their goals in very small steps. **Main Results:** We present a novel method to perform collision-free global navigation in crowded or challenging environments. Our approach is limited to scenarios with static obstacles or dynamic obstacles whose trajectories are known beforehand. The key idea is to compute bridges in the narrow passages or challenging areas of the workspace, which are collision-free regions of the workspace and have certain geometric navigation characteristics. As an agent approaches the bridge, we use the precomputed local trajectories associated with that bridge to guide an agent towards its goal position. We also present an efficient scheduling scheme which enables multiple agents to share a single bridge efficiently. The overall trajectory of each agent is calculated by using an optimal trajectory generation algorithm [3] along the interpolated path in the bridge, which combines the efficiency of the decoupled methods with the precomputed interpolating bridges. Our method can also be used for multi-agent group formation for a large number of agents.

A novel component of our approach is the computation of interpolating bridges in 2D and 3D workspaces. Each bridge lies in the collision-free space, and its boundaries are calculated using kinodynamic RRT algorithms. Our approach guarantees that when an agent enters a bridge with a velocity satisfying suitable criteria, it can always compute a collision-free trajectory that lies within the bridge. Furthermore, we present an inter-trajectory scheduling scheme for multiple agents sharing a bridge for navigation that has a small runtime overhead. We present efficient algorithms to compute these bridges in 2D and 3D workspaces, to generate trajectories within bridges at runtime, and to schedule agents sharing the bridges for efficient global navigation. We highlight the performance of our method in several challenging 2D and 3D benchmarks with narrow passages. In practice, our method can handle about 50 – 100 agents on a single CPU core for both 2D and 3D scenarios. Its worst case runtime performance can increase as a quadratic function of the number of agents, though we observe linear time performance in most scenarios.

The rest of the paper is organized as follows. We give a brief survey of prior work in Section II. We introduce our notation and provide an overview of our approach

<sup>1</sup>Liang He and Dinesh Manocha are with the Department of Computer Science, the University of North Carolina at Chapel Hill

<sup>2</sup>Jia Pan is with the Department of Mechanical and Biomedical Engineering, the City University of Hong Kong.

in Section III. We describe the bridge computation and global multi-agent navigation algorithm in Section IV and Section V, respectively. We analyze our method’s properties in Section VI, and finally highlight its performance in Section VII.

## II. RELATED WORK

There has been extensive work on multi-agent motion planning. This work includes many reactive methods such as RVO [4], HRVO [5], AVO [6], and their variants. These techniques compute a feasible movement for each agent such that it can avoid other agents and obstacles in a short time horizon. Besides robotics, they are also used for crowd simulation [7]. In practice, they are used for local navigation and collision avoidance. However, in some scenarios they may not be able to avoid the inevitable collision states (ICS) [1], [2], [8], [9] in the configuration space, due to robots’ dynamical constraints or obstacles in the scenario. Other methods [9], [10], [11] provide partial solutions to these problems, but they still cannot guarantee avoidance of all ICS in the long horizon while working in a crowd scenario with narrow passages. Even for a scenario without static obstacles, it is still difficult to perform robust collision-avoidance coordination when there are a large number of agents and high agent density [12], [13], [14].

The simplest solution to the difficulty of inevitable collision states is to design suitable protocols for multi-robot coordination/interaction [15], [16]. Some other approaches precompute roadmaps or corridors in the entire workspace to achieve high-quality path planning [17], [18]. However, these methods are not complete and may compute sub-optimal trajectories. Our method also leverages pre-computed bridges to deal with the navigation challenges in narrow or crowded regions. However, the bridges used in our approach have special properties that are useful for efficient global navigation in challenging areas in the workspace.

Centralized multi-agent navigation approaches usually leverage global single-robot planning algorithms (such as PRM or RRT) to compute a roadmap or grids for the high-level coordination [19], [20]. Compared to the decentralized methods, these algorithms compute all agents’ trajectories simultaneously and thus can better handle the complex interactions among agents. These methods can also be extended to handle non-holonomic multi-agent systems (e.g., systems composed of differential-drive robots), by using local planners like RRT [21], RRT\* [22], or other algorithms that can deal with differential constraints [3]. In addition to their benefit in terms of finding feasible trajectories, the centralized algorithms can avoid deadlock cases by leveraging high-level scheduling or coordination strategies, either coupled [23], [24] or decoupled [25], [26], [27], [28]. In general, these strategies only work in theory, because they have to ignore the robot’s dynamics and assume the robots to be operating in a discrete state space. Finally, centralized multi-agent navigation algorithms are computationally expensive and can sometimes be too slow for real-world applications.

## III. GLOBAL NAVIGATION

**Problem Definition:** Our goal is to enable a group of agents to reach their individual goals in a safe and efficient manner. In this paper, we will only consider double-integrator systems with bounded acceleration and velocities, but our approach can be extended to continuous-time multi-agent systems with general linear dynamics. During the navigation, the agents should avoid collisions with static obstacles in the environment as well as with other agents.

This problem can be formally defined as follows. We take as given a set of  $n$  decision-making agents sharing an environment containing obstacles. For simplicity, we assume the geometric shape of each agent is represented as a disc of radius  $r$ , and its current position and velocity are denoted as  $\mathbf{p}$  and  $\mathbf{v}$ , respectively. We also assume kinodynamic constraints on the agent’s motion:  $\|\mathbf{v}\| \leq v_{\max}$  and  $\|\mathbf{a}\| \leq a_{\max}$ , where  $v_{\max}$  and  $a_{\max}$  are the maximum allowed velocity and acceleration for the agent’s velocity  $\mathbf{v}$  and acceleration  $\mathbf{a}$ . The task of each agent is to compute a trajectory toward its goal position  $G$  from its initial position  $I$ . The trajectory should be collision-free and also satisfy the velocity and acceleration constraints.

**Our Approach:** Given a crowded scenario as shown in Figure 1, an efficient navigation strategy for a group of agents could be as follows: in the open space outside the narrow passage, each agent can navigate according to a path computed by optimal motion planning approaches; in the narrow area, agents should maintain a line or some formation, and then pass through the narrow region in some order.

For this purpose, we allocate a set of guidance channels called *bridges* in the narrow parts of the workspace, as shown in Figure 1. Each bridge completely lies inside the collision-free subset of the configuration space and its boundary is computed using RRT-based collision-free paths. Furthermore, the entrance of bridge makes agents can access the bridge without worry about the dynamic constraint. In other words, once agents reach the entrance, they are guaranteed to access the bridge and generate trajectories. This feature directly improves the utilization of bridge. To pass through the narrow passages of the workspace, an agent will first move toward one of the bridges according to its dynamics. However, the agent may arrive at the bridge with an arbitrary velocity. Thus, the agent must first enter an area called *entrance* in front of the bridge. The entrance will gradually adjust the agent’s velocity to make sure it satisfies the requirements with respect to that bridge. Once the agent leaves the entrance and enters the bridge, it can follow a collision-free trajectory to pass through the crowded or narrow region efficiently and safely. After leaving the bridge, the agent can switch back to the local coordination strategy and move toward its goal. In this way, the bridges can be viewed as a “highway system” for the agents to efficiently travel through challenging scenarios.

While we can build one bridge for each agent, this may result in many overlaps between the bridges, and therefore

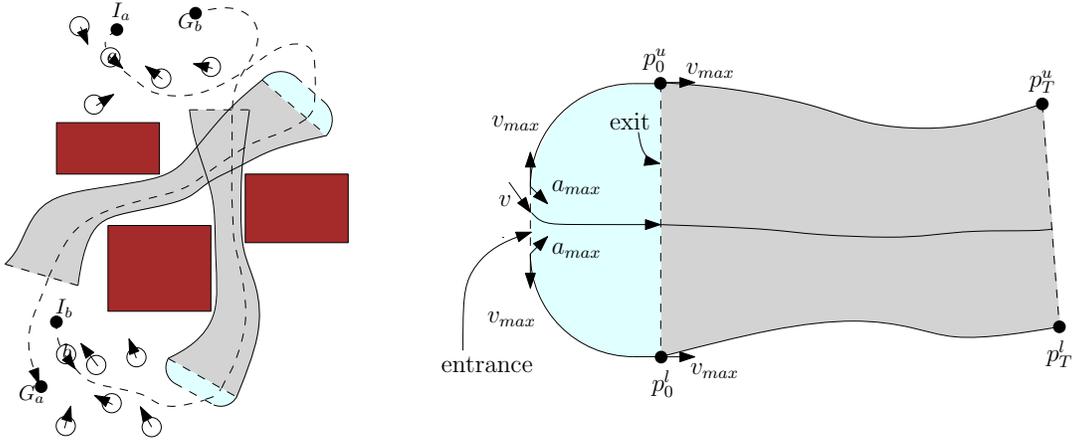


Fig. 1: **Overview of our approach:** The left figure shows a scene with three obstacles (the brown boxes), agents (the circles) want to reach their individual goals  $G_a, G_b, \dots$  from their initial positions  $I_a, I_b, \dots$ . We use two bridges (the gray regions) in the environment to help the agents navigate through the crowded areas efficiently. The agent  $a$  first goes toward the entrance (the cyan region) of one bridge following a path from  $I_a$  to the entrance. Next,  $a$  goes inside the entrance where its velocity is gradually adjusted before entering the bridge. Finally,  $a$  goes through the bridge and arrives at its goal  $G_a$ . The agent  $b$  arrives at its goal  $G_b$  through another bridge. The dashed line denotes an agent’s trajectory. These bridges enable global navigation with safety guarantees. The right figure shows the entrance to a bridge: Given an agent with velocity  $\mathbf{v}$ , we use the entrance to gradually change the agent’s velocity so that while entering the bridge the agent’s velocity is equal to the initial velocity of the bridge’s boundary trajectories.

the agents will need to slow down or even wait while moving along the bridges. Our solution is to compute a few bridges that can be reached by all agents, and allow multiple agents to share one bridge by using suitable scheduling schemes.

#### IV. BRIDGE CONSTRUCTION

In this section, we describe the details of how to automatically compute the bridges in the workspace. We will first describe our approach for a 2D workspace, which will later be extended to 3D workspaces.

We start from a zero-width bridge, which is a collision-free trajectory connecting a pair of start and goal positions in the crowded region. The trajectory is computed using a kinodynamic RRT planner [21], and has an initial velocity of  $\mathbf{v}_0$  with a magnitude  $v_{\max}$ . We then incrementally enlarge the bridge’s width until it touches one obstacle in the scenario. The built bridge has an advantageous property that as long as an agent enters the bridge with the velocity  $\mathbf{v}_0$ , it can always use an efficient interpolation scheme to calculate a collision-free trajectory passing through the bridge. However, an agent may enter a bridge with an arbitrary velocity and thus violate the bridge’s constraint on the entering velocity. To solve this problem, we add an entrance region in front of the bridge, which provides sufficient room for an agent to gradually adjust its velocity toward  $\mathbf{v}_0$  in order to leverage the bridge for safe and efficient navigation. Examples of 2D bridges and their entrances are shown in Figures 1, 2. Finally, for each entrance we compute a backward-reachable set, i.e., the set of positions from which an agent can reach the entrance. We also compute a forward-reachable set for the end line of the bridge, i.e., the set of positions that can be reached by an agent coming out of the bridge. All agents with their start positions inside the backward-reachable set and goal positions inside the forward-reachable set will leverage the

constructed bridge for navigation. We repeatedly add more bridges for the remaining agents until each agent has one associated bridge.

##### A. Iterative Bridge Enlargement

We first use the kinodynamic RRT planner [21] to connect a pair of starting and goal positions in the crowded scenarios, and the result is a trajectory  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$ . We consider the trajectory as a bridge with zero width, and its two boundary trajectories are  $\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\}$  and  $\{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\}$  respectively, where  $\mathbf{p}_i = \mathbf{p}_i^u = \mathbf{p}_i^l$  are overlapped points. Starting from this initial bridge, we incrementally enlarge the bridge’s width until it hits one obstacle. The algorithm is as shown in Algorithm 1.

##### B. Trajectory Generation in a Bridge

Given a 2D bridge, as shown in Figure 2, we can show that if the agent enters the bridge with a velocity equal to the initial velocities of the bridge’s two boundary trajectories, an efficient interpolation scheme is sufficient to generate a safe agent trajectory staying inside the bridge.

The interpolation details are described in Algorithm 2. In particular, we choose the acceleration of the agent at each time step  $i\Delta t$  as a linear interpolation of the accelerations of the corresponding waypoints on the boundary trajectories:  $\mathbf{a}_i = (1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l$ . The interpolation coefficient  $r$  is the ratio based on which the agent’s initial position  $\mathbf{p}_0$  partitions the bridge area’s start line  $\mathbf{p}_0^u\mathbf{p}_0^l$ :  $r = \frac{|\mathbf{p}_0^u\mathbf{p}_0|}{|\mathbf{p}_0^u\mathbf{p}_0^l|}$ . We can prove that in this way, the generated trajectory  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$  always stays inside the bridge:

*Theorem 1: The interpolated trajectory  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$  always stays inside the bridge.*

*Proof:* We first show that the interpolated trajectory has the following two properties  $\mathbf{p}_i = (1-r)\mathbf{p}_i^u + r\mathbf{p}_i^l$  and

---

**Algorithm 1: 2D bridge Computation Algorithm**


---

```

input : Start and goal points  $\mathbf{p}_0$  and  $\mathbf{p}_T$ 
output: A valid bridge
/* Initialize with a zero-width bridge */
1  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\} \leftarrow RRT(\mathbf{p}_0, \mathbf{p}_T)$ 
/* Set the overlapped bridge upper and lower boundaries */
2  $\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\} \leftarrow \{\mathbf{p}_0, \dots, \mathbf{p}_T\}$ 
3  $\{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\} \leftarrow \{\mathbf{p}_0, \dots, \mathbf{p}_T\}$ 
4 while true do
    /* Check collision for the bridge area between two boundaries */
5  $\text{collision} \leftarrow \text{bridgeCollide}(\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\}, \{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\})$ 
6 if collision then
7     break
    /* Change the start and goal positions of two boundaries */
8  $\mathbf{p}_0^u \leftarrow \mathbf{p}_0^u + \Delta\mathbf{p}_0, \mathbf{p}_T^u \leftarrow \mathbf{p}_T^u + \Delta\mathbf{p}_T$ 
9  $\mathbf{p}_0^l \leftarrow \mathbf{p}_0^l - \Delta\mathbf{p}_0, \mathbf{p}_T^l \leftarrow \mathbf{p}_T^l - \Delta\mathbf{p}_T$ 
    /* Generate new bridge boundaries using RRT */
10  $\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\} \leftarrow RRT(\mathbf{p}_0^u, \mathbf{p}_T^u)$ 
11  $\{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\} \leftarrow RRT(\mathbf{p}_0^l, \mathbf{p}_T^l)$ 
12 return A bridge  $b$  with boundaries  $\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\}$  and  $\{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\}$ 

```

---

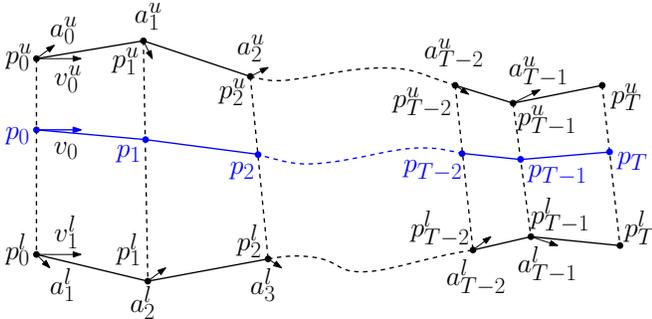


Fig. 2: Trajectory interpolation in a 2D bridge: The bridge is bounded by an upper trajectory  $\{\mathbf{p}_0^u, \mathbf{p}_1^u, \dots, \mathbf{p}_{T-1}^u, \mathbf{p}_T^u\}$  and a lower trajectory  $\{\mathbf{p}_0^l, \mathbf{p}_1^l, \dots, \mathbf{p}_{T-1}^l, \mathbf{p}_T^l\}$ , where the velocity and acceleration at each trajectory point  $\mathbf{p}_i$  are  $\mathbf{v}_i$  and  $\mathbf{a}_i$  respectively. The start line and end line of the bridge are  $\mathbf{p}_0^u \mathbf{p}_0^l$  and  $\mathbf{p}_T^u \mathbf{p}_T^l$  respectively. The agent enters the bridge at the position  $\mathbf{p}_0$  on the start line with a velocity  $\mathbf{v}_0$ . If  $\mathbf{v}_0 = \mathbf{v}_0^u = \mathbf{v}_0^l$ , the trajectory interpolation scheme can compute the agent's trajectory as  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$  which will be located completely inside the bridge.

$\mathbf{v}_i = (1-r)\mathbf{v}_i^u + r\mathbf{v}_i^l$ , for  $i = 1, \dots, T$ . We can prove these two statements by induction on  $i$ . If  $i = 1$ ,  $\mathbf{p}_0 = (1-r)\mathbf{p}_0^u + r\mathbf{p}_0^l$  is trivial because this is the definition of  $r$ ; and since  $\mathbf{v}_0 = \mathbf{v}_0^u = \mathbf{v}_0^l$  as required by the bridge's definition,  $\mathbf{v}_0 = (1-r)\mathbf{v}_0^u + r\mathbf{v}_0^l$  is also obvious. Now consider the case in which  $i > 1$ :

$$\begin{aligned}
 \mathbf{v}_i &= \mathbf{v}_{i-1} + \mathbf{a}_{i-1}\Delta t \\
 &= (1-r)\mathbf{v}_{i-1}^u + r\mathbf{v}_{i-1}^l + [(1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l]\Delta t \\
 &= (1-r)[\mathbf{v}_{i-1}^u + \mathbf{a}_i^u\Delta t] + r[\mathbf{v}_{i-1}^l + \mathbf{a}_i^l\Delta t] \quad (1) \\
 &= (1-r)\mathbf{v}_i^u + r\mathbf{v}_i^l.
 \end{aligned}$$

Similarly, we have

$$\begin{aligned}
 \mathbf{p}_i &= \mathbf{p}_{i-1} + \mathbf{v}_{i-1}\Delta t + \frac{1}{2}\mathbf{a}_{i-1}(\Delta t)^2 \\
 &= (1-r)\mathbf{p}_{i-1}^u + r\mathbf{p}_{i-1}^l + [(1-r)\mathbf{v}_{i-1}^u + r\mathbf{v}_{i-1}^l]\Delta t \\
 &\quad + \frac{1}{2}[(1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l](\Delta t)^2 \\
 &= (1-r)[\mathbf{p}_{i-1}^u + \mathbf{v}_{i-1}^u\Delta t + \frac{1}{2}\mathbf{a}_{i-1}^u(\Delta t)^2] \\
 &\quad + r[\mathbf{p}_{i-1}^l + \mathbf{v}_{i-1}^l\Delta t + \frac{1}{2}\mathbf{a}_{i-1}^l(\Delta t)^2] \\
 &= (1-r)\mathbf{p}_i^u + r\mathbf{p}_i^l. \quad (2)
 \end{aligned}$$

Based on the induction hypothesis, the waypoints of the interpolated trajectory will always be a linear interpolation of the corresponding waypoints of the bridge's two boundary trajectories. Thus, the interpolated trajectory must be contained inside the bridge. ■

---

**Algorithm 2: Trajectory generation in a 2D bridge**


---

```

input : The bridge's two boundary trajectories  $\{\mathbf{p}_0^u, \dots, \mathbf{p}_T^u\}$ 
and  $\{\mathbf{p}_0^l, \dots, \mathbf{p}_T^l\}$ , along with each waypoint  $\mathbf{p}_i$ 's
velocity  $\mathbf{v}_i$  and acceleration  $\mathbf{a}_i$ . The initial position
 $\mathbf{p}_0$  and velocity  $\mathbf{v}_0$  when the agent enters the bridge.
output: The agent's trajectory  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$ 

```

```

/* Compute the ratio into which  $\mathbf{p}_0$ 
partitions the start line  $\mathbf{p}_0^u \mathbf{p}_0^l$  */
1  $r = \frac{|\mathbf{p}_0^u \mathbf{p}_0|}{|\mathbf{p}_0^u \mathbf{p}_0^l|}$ ;
2 for  $i = 1, \dots, T-1$  do
3      $\mathbf{a}_i = (1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l$ ;
4      $\mathbf{p}_{i+1} = \mathbf{p}_i + \mathbf{v}_i\Delta t + \frac{1}{2}\mathbf{a}_i(\Delta t)^2$ ;
5      $\mathbf{v}_{i+1} = \mathbf{v}_i + \mathbf{a}_i\Delta t$ ;

```

---

### C. 2D Entrance Construction and Trajectory Generation

When an agent enters the bridge, its velocity must be the same as the initial velocity of the bridge's boundary trajectories. However, agents may arrive at the bridge with an arbitrary velocity, and thus the bridge may not be able to generate a trajectory for an agent that is completely inside the bridge. Our solution is to leave some space in front of the bridge called the *entrance* to the bridge. This space works like a buffer zone where the agent can gradually adjust its velocity to meet the bridge's requirement.

An example of the 2D entrance is shown on the right side in Figure 1, which includes four parts: one start line, one end line, and two boundary curves. The end line of the entrance is also the start line of the bridge. Each boundary curve is composed of a parabola followed by a line. The two parabolas correspond to the trajectories in which the agent's velocity is gradually rotated by 90 degrees from the initial  $\pm v_{\max}\hat{\mathbf{y}}$  to the final  $v_{\max}\hat{\mathbf{x}}$ . The velocity change is achieved by using a constant acceleration  $\mathbf{a}$ , the magnitude of which is  $a_{\max}$  and the direction of which is  $\frac{\sqrt{2}}{2}\hat{\mathbf{x}} \mp \frac{\sqrt{2}}{2}\hat{\mathbf{y}}$ . For convenience, we assume the bridge's start line is along the  $\hat{\mathbf{y}}$  direction. The two lines following the parabolas have a length  $(\sqrt{\frac{2}{3}} - \sqrt{\frac{1}{2}})\frac{v_{\max}^2}{a_{\max}}$  each.

The shape of the entrance chosen has the property that, for any agent with an arbitrary velocity  $\mathbf{v}$  ( $|\mathbf{v}| \leq v_{\max}$ ) entering the entrance, we can always find a sequence of accelerations to gradually change its velocity from  $\mathbf{v}$  to  $v_{\max}\hat{\mathbf{x}}$  (i.e., with a magnitude of  $v_{\max}$  and perpendicular to the bridge's start line). After this velocity adjustment, the agent will achieve a velocity equal to the initial velocity of the bridge's boundary trajectories when entering the bridge. According to Section IV-B, this property is required for the trajectory interpolation inside the bridge.

We can choose the acceleration such that, during the adjustment of the velocity, the agent will always stay inside the area of the entrance and is therefore, guaranteed to avoid collisions with other agents and obstacles. In particular, given an agent entering the entrance at a velocity  $\mathbf{v} = (v_x, v_y)$ , our goal is to increase  $v_x$  to  $v_{\max}$  and decrease  $v_y$  to 0. We first determine a suitable acceleration  $(a_x, a_y)$  by comparing the speed gaps in both directions. This acceleration guarantees that  $v_x$  will increase to  $v_{\max}$  after  $v_y$  arrives at 0, which is necessary to keep the agent inside the entrance region. After  $v_y$  becomes 0, the acceleration along the  $\hat{\mathbf{y}}$  will become zero, while the acceleration along  $\hat{\mathbf{x}}$  remains the same. The agent moves on until its velocity is  $v_{\max}\hat{\mathbf{x}}$ . After that, if the agent has not reached the exit line, it will continue with the current velocity. After computing the extreme positions at which the resulting trajectory will arrive, we can show that the trajectory will never go out of the entrance region and thus is guaranteed to be collision-free with static obstacles.

#### D. Bridges' Forward and Backward Reachable Regions

After generating the bridge and its entrance, we then compute its forward and backward reachable regions. We denote  $\mathcal{R}^+[\mathbf{p}, \tau]$  as the forward-reachable set of a position  $\mathbf{p}$ , (i.e., the set of positions that can be reached from  $\mathbf{p}$  with time less than  $\tau$ ), and denote  $\mathcal{R}^-[\mathbf{p}, \tau]$  as the backward-reachable set, (i.e., the set of positions that can reach  $\mathbf{p}$  with time shorter than  $\tau$ ):

$$\mathcal{R}^+(\mathbf{p}, \tau) = \{\mathbf{p}' | \text{time}(\mathbf{p}, \mathbf{p}') < \tau\} \quad (3)$$

$$\mathcal{R}^-(\mathbf{p}, \tau) = \{\mathbf{p}' | \text{time}(\mathbf{p}', \mathbf{p}) < \tau\}, \quad (4)$$

where  $\text{time}(\mathbf{x}, \mathbf{y})$  measures the time an agent needs to move from a starting position  $\mathbf{x}$  to a goal position  $\mathbf{y}$ . Both reachability sets can be efficiently estimated using the method proposed in [3].

Leveraging the concept of forward and backward reachable sets, we can compute the forward and backward reachable regions for the entire bridge as

$$\mathcal{R}^+(\text{bridge}, \tau) = \bigcup_{\mathbf{p} \in \text{bridge end line}} \mathcal{R}^+(\mathbf{p}, \tau) \quad (5)$$

$$\mathcal{R}^-(\text{bridge}, \tau) = \bigcup_{\mathbf{p} \in \text{entrance start line}} \mathcal{R}^-(\mathbf{p}, \tau). \quad (6)$$

#### E. Bridge Assignment among Agents

We have finished describing how to build a bridge, its entrance, and the corresponding reachability regions. We will now discuss how to build a set of bridges for all the agents

in the scenario, as well as how to assign a bridge to each agent. Our solution is to first build a bridge for a single agent. Next, we check whether or not there are any other agents whose initial and goal positions are within the backward and forward reachable regions of this bridge. If so, these agents can also leverage this bridge for navigation, and thus we assign this bridge to these agents. We repeatedly add more bridges until all agents have one associated bridge. The algorithm is as shown in Algorithm 3.

---

#### Algorithm 3: Generating bridges for all agents

---

```

input : Initial and goal configurations pairs
          $IG = \{(I_i, G_i)\}_{i=0}^{n-1}$  for all  $n$  agents
output: A set of bridges  $B = \{B_j\}$ 
1  $B \leftarrow \emptyset$ 
  /* compute the bridges for all of the agents */
2 while  $IG \neq \emptyset$  do
3   Select two configurations  $(I, G)$  from the set  $IG$ 
4   Construct a 2D bridge  $b$  and its entrance using Algorithm 1
5   Compute the reachability regions  $\mathcal{R}^+(b, \tau)$  and  $\mathcal{R}^-(b, \tau)$ 
  /* compute the bridge for the agents */
6   for  $(I_i, G_i) \in IG$  do
7     if  $I_i \in \mathcal{R}^-(b, \tau)$  and  $G_i \in \mathcal{R}^+(b, \tau)$  then
8       Remove  $(I_i, G_i)$  from  $IG$ 
9       Assign bridge  $b$  to the agent  $i$ 
10      Add bridge  $b$  to  $B$ 
11 return  $B$ 

```

---

#### F. 3D Bridge and Entrance

The bridge and entrance algorithms described above can be extended to the 3D workspace. Figure 3 illustrates the 3D bridge and the corresponding trajectory interpolation algorithm. The boundary of the 3D bridge is composed of  $K$  trajectories  $\{\mathbf{p}_0^k, \mathbf{p}_1^k, \dots, \mathbf{p}_{T-1}^k, \mathbf{p}_T^k\}$ , computed by the RRT algorithm, where  $k = 1, \dots, K$ . The initial points  $\mathbf{p}_0^k$  of all these  $K$  trajectories are located on the same plane (called the start plane of the bridge), and their initial velocities  $\mathbf{v}_0^k$  are the same, all having the magnitude of  $v_{\max}$ . Given an agent entering the bridge at position  $\mathbf{p}_0$  on the start plane, we can compute its trajectory using interpolation as follows. First, from all the initial points, we can select three points whose convex combination can be used for the point  $\mathbf{p}_0$ . W.l.o.g., we assume that these three points are  $\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^2$ , and their convex combination is  $\mathbf{p}_0 = u\mathbf{p}_0^0 + v\mathbf{p}_0^1 + w\mathbf{p}_0^2$ , where  $0 \leq u, v, w \leq 1$  and  $u + v + w = 1$ . Similar to Theorem 1, we can show that if we choose the acceleration at time  $t$  to be  $\mathbf{a}_t = u\mathbf{a}_t^0 + v\mathbf{a}_t^1 + w\mathbf{a}_t^2$ , the resulting trajectory will have velocity  $\mathbf{v}_t = u\mathbf{v}_t^0 + v\mathbf{v}_t^1 + w\mathbf{v}_t^2$  and position  $\mathbf{p}_t = u\mathbf{p}_t^0 + v\mathbf{p}_t^1 + w\mathbf{p}_t^2$  at time  $t$ . It follows that this trajectory will always stay inside the 3D bridge. The 3D entrance construction and trajectory generation algorithms are also similar to the 2D case, except that the entrance's boundary is now a surface bounded by a series of  $K$  parabolas.

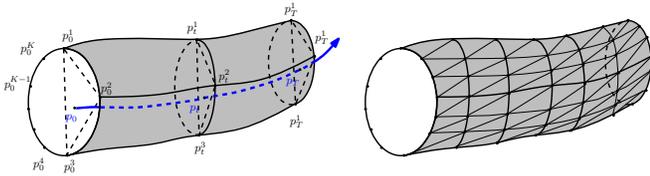


Fig. 3: The left figure shows the trajectory interpolation in a 3D bridge: the bridge is bounded by a set of boundary trajectories  $\{\mathbf{p}_0^k, \mathbf{p}_1^k, \dots, \mathbf{p}_{T-1}^k, \mathbf{p}_T^k\}$ , where  $k = 1, \dots, K$ . The agent enters the bridge at the position  $\mathbf{p}_0$  on the start plane. The trajectory interpolation is performed by first choosing three points from  $\{\mathbf{p}_i^k\}$  that can provide a convex combination for the point  $\mathbf{p}_0$ . Here we assume the first three points are chosen, and thus  $\mathbf{p}_0 = u\mathbf{p}_0^1 + v\mathbf{p}_0^2 + w\mathbf{p}_0^3$ , where  $0 \leq u, v, w \leq 1$  and  $u + v + w = 1$ . Similar to the 2D bridge, we can always choose a suitable acceleration  $\mathbf{a}_i$  so that the resulting trajectory  $\{\mathbf{p}_0, \dots, \mathbf{p}_T\}$  completely locates inside the bridge. The right figure shows the triangulation for the 3D bridge shown in the left figure. In this case, each point is one waypoint in one of the  $K$  boundary trajectories.

### G. Collision Checking for a Bridge

Given a bridge, we need to make sure that it is collision-free. To perform collision checking on a bridge, we triangulate the bridge boundary, and then perform collision checking between the triangulated bridge boundary with static obstacles in the environment. The right side in figure 3 shows the triangulation result for the 3D bridge. The collision checking is performed using traditional bounding volume techniques [29].

## V. GLOBAL NAVIGATION USING BRIDGES

Once the bridges and their entrances are computed, we can leverage them for efficient multi-agent global navigation. Each agent will first move toward the bridge assigned to it along an optimal trajectory computed using [3]. Once it reaches the entrance, it can enter the entrance and then go through the crowded or narrow area by following the trajectory interpolated by the bridge. After leaving the bridge, the agent switches back to moving toward its individual goal following an optimal trajectory computed using [3].

### A. Inter-Trajectory Scheduling

However, there is still one unresolved situation: several agents may try to leverage the bridge at the same time, and this may result in collisions between the agents inside the bridge. To avoid this problem, we use a scheduling scheme among the agents so that they can share the bridge in a safe and efficient manner for collision-free navigation. This is achieved by inter-trajectory scheduling.

Algorithm 4 shows a simple scheduling algorithm among trajectories. In particular, we check whether a planned trajectory will collide with any previously scheduled trajectory. If so, we delay the trajectory for a small time  $\delta t$ . This process continues until all collisions among trajectories are resolved.

## VI. ASYMPTOTIC ANALYSIS

The complexity of our method's online phase is  $\mathcal{O}(n^2)$ , where  $n$  is the number of agents. In particular, the inter-trajectory scheduling in Section V-A dominates the runtime

---

### Algorithm 4: Inter-bridge scheduling: used for collision-free trajectory computation

---

```

input : Original plans  $P = \{P_i\}_{i=0}^n$ 
output: Scheduled plans  $P' = \{P'_i\}_{i=0}^n$ 
1 for  $i = 1$  to  $n$  do
2    $P'_i \leftarrow P_i$ 
3   for  $j = i - 1$  to  $0$  do
4     while  $P'_i$  collide with  $P'_j$  do
5        $P'_i$  postpone  $\delta t$ 
6 return  $P'$ 

```

---

cost. As shown in line 5 of Algorithm 4, each agent needs to check whether its planned trajectory will collide with any plans of previously scheduled agents. If so, the agent will delay its plan to avoid collisions. For  $n$  agents, we need to execute  $\frac{1}{2}n^2$  checks in total. For each check between two trajectories, we need to further check whether any two segments from these paths collide with each other. For two long trajectories  $P$  and  $P'$ , we may need to perform a pairwise checking that can have  $|P| \times |P'|$  complexity. However, note that when agents are moving inside the same bridge, their speeds are determined by their position in the start line and are fixed while passing through the bridge (Theorem 1). As a result, agents falling behind will never catch up the agents in the front, and thus we can determine the collision status between two plans by only checking collisions between a few segments. In fact, we observe that at most  $C < 10$  collision checks are performed between two trajectories in all our benchmarks. In this way, the worst case computational complexity of the online phase is  $C \times \frac{1}{2}n^2 = \mathcal{O}(n^2)$ . In practice, the running time is close to  $\mathcal{O}(n)$  [30] [31].

## VII. EXPERIMENTS

In this section, we demonstrate the time cost of our navigation method on four challenging benchmarks, including two 3D scenarios and two 2D scenarios as shown in Figure 4. The results for the time cost are shown in Table I. We break down the time costs into the offline precomputation and the online execution steps. In the offline stage, we precompute a set of bridges in about 200-1000 ms. In the online step, the time cost includes three parts: the time to compute the trajectories outside bridges, the time to compute the trajectories inside bridges, and the time to scheduling trajectories of different agents. The overheads to calculate trajectories inside and outside bridges are both quite small, less than 1 ms and 0.5s respectively. The scheduling algorithm takes about 1s on all benchmarks. We also compare the performance of our method with two other state-of-the-art local navigation approaches: one uses the Reciprocal Velocity Obstacle (RVO) techniques ([4] for 3D benchmarks and [5] for 2D benchmarks), and the other is a combination of RVO and the global planner PRM (RVO+PRM). We compare the simulation time required by different methods to make all agents reach the goal, and the number of collisions

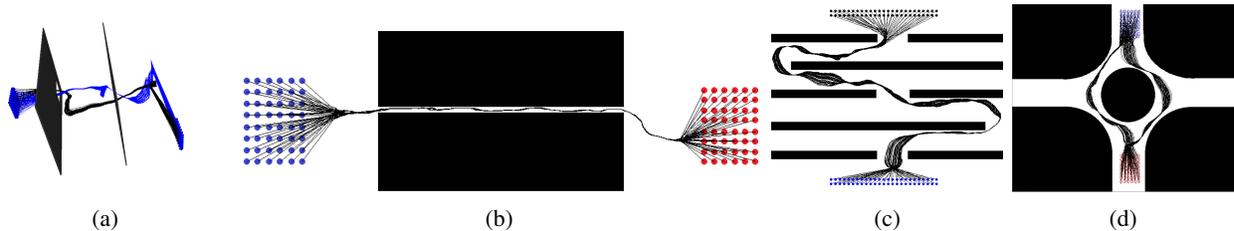


Fig. 4: We use four challenging 2D and 3D benchmarks with narrow passages or crowded regions to evaluate the performance of our method. The collision-free paths computed by our method are shown as the black curves with the blue and red points as start and goal positions, respectively.

occurred during the simulation. The comparison results are shown in Table II. We provide more experimental details in our technical report [30]. We implemented all experiments using C++ on an Intel Core i7 CPU running at 3.30GHz with 16GB of RAM.

*a) Benchmark 1:* This is a challenging scenario where 96 agents need to pass through two holes simultaneously and arrive at their respective goals, as shown in Figure 4(a). The holes are small and only allow two agents to pass through at the same time. However, the areas around the agents' start and goal positions are widely open, and thus the agents can easily find a path to enter the bridges, as shown by the small outside bridge cost in Table I. Compared to RVO algorithms, our method can effectively compute collision-free global paths, while the RVO local navigation methods get stuck in the narrow passages and can take a long time to make all agents reach the goals. In addition, our method results in fewer collision cases as compared to RVO and RVO+PRM, as shown in Table II.

*b) Benchmark 2:* This is a 2D benchmark with a long narrow corridor where 96 agents are trying to move from one end to the other end, as shown in Figure 4(b). The corridor only allows one agent to pass through at a time. For this scene, we compute two bridges connecting the two open spaces in the scenario. The RVO method can compute the collision-free trajectories for agents, but they tend to get stuck for a while due to the narrow corridor. The RVO+PRM method also results in a significant number of collisions. Our method generates a more stable and faster simulation than RVO and RVO+PRM, as shown in Figure II.

*c) Benchmark 3:* This scenario has multiple narrow passages as shown in Figure 4(c), and we construct two bridges to help the global navigation through these narrow passages. In this scenario, the RVO method fails to compute collision-free paths because agents are getting stuck in the narrow passages. The RVO+PRM method can find a feasible solution, but it takes significantly longer time as compared to our method and also results in more collisions, as demonstrated in Table II.

*d) Benchmark 4:* In this benchmark, we have a circle obstacle with narrow passages in the scene. Our method computes two bridges to connect the upper and lower regions in the workspace. The simulation result is shown in Figure 4(d). As compared to local methods, our approach results in fewer agent-agent collisions and can compute the final trajectories

faster as shown in Table II.

## VIII. CONCLUSION AND FUTURE WORK

We present a novel multi-agent global navigation algorithm using interpolation bridges. Our approach is general and overcomes some of the major limitations of prior methods in terms of navigating through crowded areas or narrow regions. We present new techniques to compute these bridges in 2D and 3D workspaces and use their properties to compute interpolating collision-free trajectories for the agents. The construction of our bridge enables collision-free multi-agent global navigation. We have demonstrated its performance on many complex 2D and 3D scenarios and can perform collision-free navigation for tens of agents in real time.

Our approach has some limitations. The bridge computation is limited to static obstacles or dynamic obstacles whose trajectories are known a priori. The complexity of global navigation increases with the number of bridges in the workspace, and very complex scenarios can result in a high number of bridges. Furthermore, our current approach is limited to agents with linear dynamics. There are many avenues for future work. In addition to overcoming these limitations, we would like to design improved algorithms for bridge computation and further evaluate their performance in complex scenarios. Furthermore, we would like to extend this approach so that it can provide completeness guarantees, similar to [32].

## IX. ACKNOWLEDGEMENT

This research is supported in part by ARO grant W911NF-14-1-0437 and NSF grant 1305286. Jia Pan is supported by HKSAR Research Grants Council (RGC) General Research Fund (GRF), CityU 17204115, 21203216, and NSFC/RGC Joint Research Scheme CityU103/16.

## REFERENCES

- [1] T. Fraichard and H. Asama, "Inevitable collision states. a step towards safer robots?" in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, 2003, pp. 388–393 vol.1.
- [2] T. Fraichard, "A short paper about motion safety," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 1140–1145.
- [3] D. Webb and J. van den Berg, "Kinodynamic rrt\*: Asymptotically optimal motion planning for robots with linear dynamics," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 5054–5061.
- [4] J. van den Berg, S. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research*, ser. Springer Tracts in Advanced Robotics, C. Pradaliere, R. Siegwart, and G. Hirzinger, Eds. Springer Berlin Heidelberg, 2011, vol. 70, pp. 3–19.

	#agents	#bridges	offline	online		
			bridge construction	scheduling	outside bridge	inside bridge
Benchmark 1	96	2	204,513	1,012	23	< 1
Benchmark 2	96	2	453,267	1,432	47	< 1
Benchmark 3	100	2	553,462	1,431	132	< 1
Benchmark 4	96	2	1079,381	1532	560	< 1

TABLE I: The running time (ms) for bridge construction and scheduling, and the average running time for each agent to compute its trajectory inside the outside the bridges on various benchmarks. All the timings are in milliseconds on a single CPU core.

	#collisions			#frames		
	RVO	RVO+PRM	ours	RVO	RVO+PRM	ours
Benchmark 1	-	13	0	-	3905	2024
Benchmark 2	212	9	0	3938	2215	1482
Benchmark 3	-	11	0	-	1101	762
Benchmark 4	517	4	0	4882	752	417

TABLE II: The comparison of the number of collisions and simulation steps between our method, the local navigation method RVO, and the combination of RVO and a global planner (RVO+PRM). Our global approach results in fewer collisions and the agents can reach the goal positions in fewer time steps.

- [5] J. Snape, S. Guy, J. van den Berg, and D. Manocha, "Smooth coordination and navigation for multiple differential-drive robots," in *Experimental Robotics*, ser. Springer Tracts in Advanced Robotics, O. Khatib, V. Kumar, and G. Sukhatme, Eds. Springer Berlin Heidelberg, 2014, vol. 79, pp. 601–613.
- [6] J. Van Den Berg, J. Snape, S. J. Guy, and D. Manocha, "Reciprocal collision avoidance with acceleration-velocity obstacles," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 3475–3482.
- [7] S. Curtis, A. Best, and D. Manocha, "Menge: A modular framework for simulating crowd movement," *Collective Dynamics*, vol. 1, pp. 1–40, 2016.
- [8] J. Reif and M. Sharir, "Motion planning in the presence of moving obstacles," in *Symposium on Foundations of Computer Science*, 1985, pp. 144–154.
- [9] K. E. Bekris, D. K. Grady, M. Moll, and L. E. Kavraki, "Safe distributed motion coordination for second-order systems with different planning cycles," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 129–150, 2012.
- [10] L. He and J. van den Berg, "Meso-scale planning for multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 2839–2844.
- [11] L. He, J. Pan, W. Wang, and D. Manocha, "Proxemic group behaviors using reciprocal multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 292–297.
- [12] K. Solovey, J. Yu, O. Zamir, and D. Halperin, "Motion planning for unlabeled discs with optimality guarantees," in *Robotics: Science and Systems*, 2015.
- [13] K. Solovey and D. Halperin, "On the hardness of unlabeled multi-robot motion planning," *CoRR*, vol. abs/1408.2260, 2014. [Online]. Available: <http://arxiv.org/abs/1408.2260>
- [14] J. Yu and S. M. LaValle, "Optimal multi-robot path planning on graphs: Complete algorithms and effective heuristics," *IEEE Transactions on Robotics*, vol. to appear, 2015.
- [15] K. Dresner and P. Stone, "A multiagent approach to autonomous intersection management," *Journal of Artificial Intelligence Research*, vol. 31, no. 1, pp. 591–656, Mar. 2008.
- [16] R. Alami, S. Fleury, M. Herrb, F. Ingrand, and F. Robert, "Multi-robot cooperation in the martha project," *IEEE Robotics Automation Magazine*, vol. 5, no. 1, pp. 36–47, 1998.
- [17] R. Geraerts, A. Kamphuis, I. Karamouzas, and M. Overmars, "Using the corridor map method for path planning for a large number of characters," in *Motion in Games*. Springer, 2008, pp. 11–22.
- [18] R. Wein, J. van den Berg, and D. Halperin, "Planning high-quality paths and corridors amidst obstacles," *The International Journal of Robotics Research*, vol. 27, no. 11-12, pp. 1213–1231, 2008.
- [19] J. van den Berg and M. Overmars, "Roadmap-based motion planning in dynamic environments," *IEEE Transactions on Robotics*, vol. 21, no. 5, pp. 885–897, 2005.
- [20] —, "Prioritized motion planning for multiple robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2005, pp. 430–435.
- [21] S. LaValle and J. Kuffner, "Randomized kinodynamic planning," *International Journal on Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [22] S. Karaman and E. Frazzoli, "Optimal kinodynamic motion planning using incremental sampling-based methods," in *IEEE Conference on Decision and Control*, 2010, pp. 7681–7687.
- [23] M. Čáp, P. Novák, J. Vokřínek, and M. Pěchouček, "Multi-agent rrt: Sampling-based cooperative pathfinding," in *International Conference on Autonomous Agents and Multi-agent Systems*, 2013, pp. 1263–1264.
- [24] D. Ferguson, N. Kalra, and A. Stentz, "Replanning with rrt," in *IEEE International Conference on Robotics and Automation*, 2006, pp. 1243–1248.
- [25] M. Katsev, J. Yu, and S. LaValle, "Efficient formation path planning on large graphs," in *IEEE International Conference on Robotics and Automation*, 2013, pp. 3606–3611.
- [26] G. Sanchez and J.-C. Latombe, "Using a prm planner to compare centralized and decoupled planning for multi-robot systems," in *IEEE International Conference on Robotics and Automation*, vol. 2, 2002, pp. 2112–2119 vol.2.
- [27] J. van den Berg, J. Snoeyink, M. Lin, and D. Manocha, "Centralized path planning for multiple robots: Optimal decoupling into sequential plans," in *Robotics: Science and systems*, 2009.
- [28] R. Luna and K. E. Bekris, "Push and swap: Fast cooperative pathfinding with completeness guarantees," in *Joint Conference on Artificial Intelligence - Volume Volume One*, 2011, pp. 294–300.
- [29] S. Gottschalk, M. C. Lin, and D. Manocha, "Obbtrees: A hierarchical structure for rapid interference detection," in *SIGGRAPH*, 1996, pp. 171–180.
- [30] L. He, J. Pan, and D. Manocha, "Efficient multi-agent global navigation using interpolating bridges," 2016, technical Report. <http://gamma.cs.unc.edu/Proxemic/multibody.pdf>.
- [31] M. Čáp, J. Gregoire, and E. Frazzoli, "Provably safe and deadlock-free execution of multi-robot plans under delaying disturbances," *CoRR*, vol. abs/1603.08582, 2016. [Online]. Available: <http://arxiv.org/abs/1603.08582>
- [32] L. He, B. Jia, Z. Pan, and D. Manocha, "Resolution-complete multi-agent motion planning with arbitrarily-shaped obstacles," University of North Carolina at Chapel Hill, Tech. Rep., 2016.