

Near-Optimal Multi-Robot Motion Planning with Finite Sampling

Dror Dayan¹, Kiril Solovey², Marco Pavone², and Dan Halperin¹

Abstract—An underlying structure in several sampling-based methods for continuous multi-robot motion planning (MRMP) is the tensor roadmap (TR), which emerges from combining multiple PRM graphs constructed for the individual robots via a tensor product. We study the conditions under which the TR encodes a near-optimal solution for MRMP—satisfying these conditions implies near optimality for a variety of popular planners, including dRRT*, and the discrete methods M* and CBS when applied to the continuous domain. We develop the first finite-sample analysis of this kind, which specifies the number of samples, their deterministic distribution, and magnitude of the connection radii that should be used by each individual PRM graph, to guarantee near-optimality using the TR. This significantly improves upon a previous asymptotic analysis, wherein the number of samples tends to infinity. Our new finite sample-size analysis supports guaranteed high-quality solutions in practice within finite time. To achieve our new result, we first develop a sampling scheme, which we call the *staggered grid*, for finite-sample motion planning for individual robots, which requires significantly less samples than previous work. We then extend it to the much more involved MRMP setting which requires to account for interactions among multiple robots. Finally, we report on a few experiments that serve as a verification of our theoretical findings and raise interesting questions for further investigation.

I. INTRODUCTION

Multi-robot (MR) systems are already playing a crucial role in manufacturing, warehouse automation, and natural resource monitoring, and in the future they will be employed in even broader domains from space exploration to search-and-rescue. One of the most basic ingredients necessary in all those applications are mechanisms for multi-robot motion planning (MRMP), which should quickly generate motion trajectories to move robots from their origins to destinations, while avoiding collisions with the environment and between robots. In many cases, it is desirable to develop MRMP approaches that provide strong guarantees of completeness and near-optimality, to ensure that a high-quality solution would be found (if one exists). To achieve this, methods for MRMP must accurately capture the continuous state space of individual robots and the intricate interactions between multiple robots. Those considerations make the task of designing efficient high-quality methods for MRMP tremendously challenging [1–4].

In this work we develop a general approach for centralized near-optimal sampling-based motion planning using a finite number of samples, which is the first of its kind. Previous

solutions guarantee path quality for the multi-robot case only asymptotically, namely, as the number of samples tends to infinity. In contrast, our new method explicitly prescribes finite samples sets, and by that supports guaranteed high-quality solutions in practice, within bounded running time.

Related work. A common approach to MRMP, which is often taken in the AI research community, is to consider a discretized version of the problem, termed multi-agent pathfinding (MAPF), wherein robots are assumed to move along vertices of a graph. A variety of methods were developed for MAPF, including integer-programming formulations [5], path-based search methods [6], and conflict-based search [7]. Those have been successfully applied to a variety of problems involving multiple robots—from warehouse management [8] to multi-drone package delivery [9]. Unfortunately, MAPF methods provide no solution quality guarantees with respect to the original continuous MRMP problem, since they usually employ a crude discretization of the robots’ environment in the form of a regular lattice, where a cell size corresponds to a robot’s bounding box.

Recent work extends conflict-based search to the continuous MRMP domain, by using PRM graphs to capture the individual robots’ state space [10–12]. Nevertheless, no guidelines are provided as to how to construct those roadmaps (in terms of number of samples, their distribution, and connection radius) to guarantee completeness or optimality with respect to the original MRMP problem.

In a different line of work, computational geometry methods are employed to explicitly reason about the robots’ continuous state space. Such methods are quite powerful, in that they guarantee polynomial runtime, completeness, and near-optimal solutions for different quality metrics [13–16]. However, those methods are typically restricted to disc-shaped robots operating within a planar domain, and they require special separation constraints, e.g., between the robots’ initial and terminal positions, in order to work correctly, which limits their applicability in practice.

A promising direction, which aims to overcome the limitations of discrete and geometry-based approaches, are sampling-based (SB) planners. SB-planners were initially developed to tackle the single-robot motion-planning problem for complex systems, by using random sampling of states to capture the structure of the robot’s complex state space, which results in a discrete graph representation. Many of those algorithms, including the celebrated PRM [17–19], and RRT [20, 21] algorithms are known to converge to a solution asymptotically with the number of samples drawn. Some SB-planners are also guaranteed to converge to the optimal solution as the number of samples tends to infinity [22–26].

Consequently, a variety of SB-MRMP approaches have emerged, from methods that aim to apply PRM-based so-

¹ D. Dayan and D. Halperin are with the Blavatnik School of Computer Science, Tel-Aviv University, Israel. Their work was supported in part by the Israel Science Foundation (grant no. 1736/19), by the US NSF/US-Israel BSF (grant no. 2019754), by the Israel Ministry of Science and Technology (grant no. 103129), by the Blavatnik Computer Science Research Fund, and by a grant from Yandex.

² K. Solovey and M. Pavone are with the Aeronautics and Astronautics Department, Stanford University, CA 94305, USA. Their work was supported in part by the Toyota Research Institute (TRI) and the Center for Automotive Research at Stanford (CARS).

lutions [27, 28], to techniques that sample local instances of the MAPF problem [29, 30]. A different method, termed discrete RRT (dRRT) [31], carefully explores an implicitly-represented tensor roadmap (TR), which emerges from combining several PRM roadmaps constructed for the individual robots, to effectively solve instances of MRMP requiring tight coordination between multiple robots. A recent work further improves this approach with the dRRT* method [32], which is also shown to yield a near-optimal MRMP solution. To the best of our knowledge, this is the only scalable SB-MR planner with such a guarantee. Unfortunately, this result is asymptotic and does not specify guarantees for a finite number of samples.

Contribution. We develop a general framework for near-optimal SB-MRMP while using a finite number of samples. To this end, we study the structure of the TR which is an underlying ingredient in dRRT*, MC-CBS [11] (a continuous-space extension of CBS), and M* [6] (when applied to a continuous space). We develop conditions under which the TR encodes a near-optimal solution to MRMP—if those conditions hold, it implies that the aforementioned planners are guaranteed to be near-optimal as well. In particular, we prescribe a recipe for constructing individual-robot PRM graphs, in terms of number of samples, their deterministic distribution, and connection radius, so that the resulting TR encodes a near-optimal solution for MRMP.

To achieve this we refine the asymptotic analysis that we developed in [32] (in the context of the dRRT* algorithm) to the finite-sample regime. The latter requires a much more careful study of the different ingredients of the problem, including the clearance parameter between robots and obstacles, and the approximation factor, which were previously assumed to be infinitesimally small. A key component in our result is a new sampling scheme that we develop, which we call the *staggered grid*, for finite-sample motion planning for individual robots. This sampling scheme requires significantly less samples than our previous work [33], to achieve near-optimality for the single-robot case.

The organization of this paper is as follows. In Section II we describe the single-robot problem, introduce the staggered grid, and study the theoretical properties of PRM using this sampling scheme. In Section III we describe our central contribution, namely the extension of this theoretical result to the multi-robot setting. We provide experimental results in Section IV, and conclude with an outline of future work in Section V. An extended version of this paper provides additional information and missing proofs [34].

II. IMPROVED SAMPLING DISTRIBUTIONS FOR A SINGLE ROBOT

We present improved sampling distributions for probabilistic roadmaps (PRM) for the *single-robot* case. First, we introduce basic ingredients of the problem, then review the concept of PRM, and introduce our staggered-grid sampling. We conclude with our main theoretical result for the single-robot case, pertaining to the quality of solutions obtained using PRM with staggered-grid sampling.

A. Basics of single-robot motion planning

Let $\mathcal{C}$ denote the configuration space of the robot, which we assume here to be $[0, 1]^d$, where d is the dimension,

which is the number of degrees of freedom of the robot. The free space, denoted by $\mathcal{C}^f \subset \mathcal{C}$, represents the set of all collision-free configurations. A motion planning problem is a tuple $\mathcal{M} := (\mathcal{C}^f, x^s, x^g)$, where $\mathcal{C}^f$ is the free space, and $x^s, x^g \in \mathcal{C}^f$ are the start and goal configurations, respectively. A solution of $\mathcal{M}$ is a continuous collision-free trajectory $\sigma : [0, 1] \rightarrow \mathcal{C}^f$ that begins at $\sigma(0) = x^s$ and ends at $\sigma(1) = x^g$. We measure the quality of a trajectory σ by its length, which is denoted by $\|\sigma\|$.

A crucial property of trajectories in sampling-based planning is the notion of clearance. A trajectory σ has δ -clearance if $\bigcup_{0 \leq t \leq 1} B_\delta(\sigma_i(t)) \subseteq \mathcal{C}^f$, for $B_\delta(\sigma_i(t))$ being the d -dimensional closed Euclidean ball with radius δ centered at $\sigma_i(t)$. We say that $\mathcal{M}$ is δ -clear if there exists a trajectory σ with clearance δ that solves $\mathcal{M}$.

B. Probabilistic roadmaps and sample sets

We provide a formal definition of the Probabilistic Roadmap (PRM) method [17], which constructs a discrete graph that captures the connectivity of $\mathcal{C}^f$ via sampling. PRM plays a critical role in various sampling-based planners (see, e.g., [33, 35]). PRM is also instrumental to our result both on single-robot motion-planning in this section and on multi-robot motion-planning in Section III.

For a given motion-planning problem $\mathcal{M} = (\mathcal{C}^f, x^s, x^g)$, a sample (point) set $\mathcal{X} \subset \mathcal{C}^f$, and a connection radius $r > 0$, PRM generates a graph denoted by $G_{\mathcal{M}(\mathcal{X}, r)} = (V, E)$. The vertex set V consists of all the collision-free configurations in $\mathcal{X} \cup \{x^s, x^g\}$. The set of (undirected) edges, E , consists of all vertex pairs $v, u \in V$ such that the Euclidean distance between them is at most r , and the straight-line segment between them is collision-free. Formally, we define

$$V := (\mathcal{X} \cup \{x^s, x^g\}) \cap \mathcal{C}^f, \text{ and} \\ E := \{\{v, u\} \in V \times V : \|v - u\| \leq r, \text{CH}(\{v, u\}) \subset \mathcal{C}^f\},$$

where CH denotes the convex hull of a point set.

To measure the quality of a sample set $\mathcal{X}$ and a connection radius r we use the following definition [33].

Definition 1 (Single-robot (ε, δ) -completeness). Given a sample set $\mathcal{X}$ and connection radius r , we say that $(\mathcal{X}, r)$ is (ε, δ) -complete for some stretch $\varepsilon > 0$ and clearance parameter $\delta > 0$ if for every δ -clear $\mathcal{M} = (\mathcal{C}^f, x^s, x^g)$ it holds that

$$d(G_{\mathcal{M}(\mathcal{X}, r)}, x^s, x^g) \leq (1 + \varepsilon) \text{OPT}_\delta,$$

where $d(G_{\mathcal{M}(\mathcal{X}, r)}, x^s, x^g)$ denotes the length of the shortest trajectory from x^s to x^g in the graph $G_{\mathcal{M}(\mathcal{X}, r)}$, and OPT_δ is the length of the shortest δ -clear solution to $\mathcal{M}$.

C. Efficient sampling via a staggered grid

In preparation for our main result in this section we introduce a new sampling scheme termed a *staggered grid*, which we denote by $\mathcal{X}_{\beta, \gamma}$, and study its implications with respect to (ε, δ) -completeness of PRM. Refer to Figure 1 for an illustration. An important property of this staggered grid, which we prove in the extended version [34], is that by placing β -radius hyperspheres centered at the points $\mathcal{X}_{\beta, \gamma}$ we obtain a coverage of the robot's configuration space. Moreover, the size of the set $\mathcal{X}_{\beta, \gamma}$ is smaller than previously obtained sets with similar coverage properties [33]. Notice that we use a slightly shrunk instance of the unit hyper-cube; a solution path should not pass too close to the boundary of

the hyper-cube to respect the clearance condition—this is the role of the parameter γ .

Definition 2 (Staggered grid). For given $\beta > 0, \gamma > 0$ the staggered grid $\mathcal{X}_{\beta,\gamma}$ is a union of two point sets $\mathcal{X}_{\beta,\gamma}^1, \mathcal{X}_{\beta,\gamma}^2$ in $[0, 1]^d$, where

$$\mathcal{X}_{\beta,\gamma}^l = \left\{ (p_1, p_2, \dots, p_d) : p_i = \gamma + (2k + l - 2)w, \right. \\ \left. 2 - l \leq k \leq \left\lceil \frac{1 - 2\gamma}{2w} \right\rceil, 1 \leq i \leq d \right\},$$

for $l \in \{1, 2\}$ and $w = \beta\sqrt{2}/\sqrt{d}$.

Notice that both $\mathcal{X}_{\beta,\gamma}^1$ and $\mathcal{X}_{\beta,\gamma}^2$ are square grids of side length $2w$, and the total number of points in the construction is $\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + 1\right)^d$.

Next we make the connection between the staggered grid and a PRM graph that is (ε, δ) -complete. Namely, we will show that given ε and δ , there is a sample set and radius $(\mathcal{X}_{\beta,\gamma}, r)$ that are (ε, δ) -complete, where each of β and r depends on both ε and δ , and γ is equal to δ .

Theorem 1 (Sufficient conditions for (ε, δ) -completeness). Fix a stretch parameter $\varepsilon > 0$ and clearance $\delta > 0$. For a sampling distribution $\mathcal{X} = \mathcal{X}_{\alpha,\delta}$, where $\alpha = \frac{\varepsilon}{\sqrt{1+\varepsilon^2}}$, and the radius $r = \frac{2(\varepsilon+1)}{\sqrt{1+\varepsilon^2}}\delta$, it follows that $(\mathcal{X}, r)$ is (ε, δ) -complete.

Our proof of this theorem, which appears in the extended version [34], is a refined version of [33, Proof of Theorem 2]. In particular, we first prove that the staggered grid $\mathcal{X}_{\alpha,\delta}$ provides an $\alpha\delta$ -cover of $\mathcal{C}^f$, in the sense that for every point $c \in \mathcal{C}^f$ there exists $x \in \mathcal{X}_{\alpha,\delta}$ such that $\|x - c\| \leq \alpha\delta$. We then use this property to show that we can use the vertices and edges of $G_{\mathcal{M}(\mathcal{X}_{\alpha,\delta}, r)}$ to closely follow a shortest δ -clear path for $\mathcal{M}$. The value α is set so that the resulting solution is of cost at most $(1 + \varepsilon)\text{OPT}_\delta$.

The previous proof [33] relies on a different sampling distribution, which we denote by $\mathcal{X}_{\text{prev}}$, and the same connection radius as above. The distribution $\mathcal{X}_{\text{prev}}$ is constructed by iteratively placing samples at points that are at a distance at least $\alpha\delta$ from the nearest previous sample, which makes it difficult to construct in practice due to the need to keep track of a hypersphere arrangement [36]. In contrast, our distribution can be easily constructed from Definition 2. Moreover, our sampling distribution $\mathcal{X}_{\alpha,\delta}$, which is sufficient for (ε, δ) -completeness, is significantly smaller in terms of number of sample points. In the extended version of our paper [34], we show that $\frac{|\mathcal{X}_{\text{prev}}|}{|\mathcal{X}_{\alpha,\delta}|} \approx \frac{\sqrt{\pi d}}{2} 1.3687^d$ for small values of ε and δ .




Fig. 1: Illustration of Definition 2 in two dimensions, for $\beta = 0.08, \gamma = 0.1$. On the left we visualize the first layer $\mathcal{X}_{\beta,\gamma}^1$ of the staggered grid (denoted by green discs). On the right we add the second layer $\mathcal{X}_{\beta,\gamma}^2$ (denoted by red discs). The centers of the discs are the points of $\mathcal{X}_{\beta,\delta}$.

We also provide a comparison for specific values of δ, ε, d where we demonstrate a reduction of one order of magnitude with respect to the sample size in our favor and compare our result to the lower-bound obtained in [33, Theorem 1].

III. NEAR-OPTIMAL TENSOR ROADMAPS FOR MRMP

We present our central contribution: we extend our results from the previous section to the multi-robot setting.

A. Basics of multi-robot motion planning

We provide a definition of the multi-robot motion planning (MRMP) problem. We consider the setting of $R \geq 2$ identical robots operating in a shared workspace, and denote by $\mathcal{C}_i \subset [0, 1]^d$ the configuration space of robot i , $1 \leq i \leq R$. We define $\mathcal{C}_i^f, \mathcal{C}_i^o \subset \mathcal{C}_i$ to be the free and forbidden spaces, respectively, of robot i . Since the robots are identical it holds that $\mathcal{C}_i^f = \mathcal{C}_j^f$ for all $1 \leq i \leq j \leq R$.

The configuration space of the multi-robot system $\mathbb{C}$, termed the *composite configuration space*, is the Cartesian product of the individual robots' configuration spaces, i.e., $\mathbb{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_R$. That is, a composite configuration $Q = (q_1, \dots, q_R) \in \mathbb{C}$ is an R -tuple of single-robot configurations, where $q_i \in \mathcal{C}_i$. For two distinct robots i, j , we denote by $I_i^j(q_j) \subset \mathcal{C}_i$ the set of configurations of robot i that lead to collision with robot j when j is at configuration q_j . The composite free space $\mathbb{C}^f \subset \mathbb{C}$ consists of all composite configurations $(q_1, \dots, q_R)$ such that (i) $q_i \in \mathcal{C}_i^f$ for every $1 \leq i \leq R$, and (ii) $q_i \notin I_i^j(q_j)$ for every $1 \leq i \neq j \leq R$, which ensure that robot-obstacle and robot-robot collisions are avoided, respectively.

Given start and goal positions $x_i^s, x_i^g \in \mathcal{C}_i^f$, respectively, for each robot $1 \leq i \leq R$, define $\vec{x}^s = (x_1^s, \dots, x_R^s)$ and $\vec{x}^g = (x_1^g, \dots, x_R^g)$. The MRMP problem, denoted by $\mathbb{M} = (\mathbb{C}^f, \vec{x}^s, \vec{x}^g)$, consists of finding trajectories for the R robots such that the robots begin their motion at $\vec{x}^s$, end at $\vec{x}^g$, and avoid collisions (both with obstacles and with each other) along the way. Formally, the objective is to find a collision-free composite trajectory of the form $\Sigma : [0, 1] \rightarrow \mathbb{C}^f$, where Σ is an R -tuple $\Sigma = (\sigma_1, \dots, \sigma_R)$ of single-robot trajectories $\sigma_i : [0, 1] \rightarrow \mathcal{C}_i^f$, such that $\Sigma(0) = \vec{x}^s, \Sigma(1) = \vec{x}^g$, and $\Sigma(\tau) \in \mathbb{C}^f$ for all $0 \leq \tau \leq 1$.

In this work we are interested in finding high-quality solutions for the multi-robot problem. We consider as cost criterion the sum of single-robot trajectory lengths, denoted as $\text{cost}(\Sigma) = \sum_{i=1}^R \|\sigma_i\|$. Our analysis can be easily adapted to the case where the cost criterion is the maximum of trajectory lengths, i.e., $\max_{i=1}^R \|\sigma_i\|$ (see [34]).

B. Tensor roadmaps

We provide a formal definition of the tensor roadmap (TR), which is implicitly explored by sampling-based planners such as dRRT [31], dRRT* [32], as well as by search-based methods such as MC-CBS [11] and M* [6].

For every robot i , $1 \leq i \leq R$, let $G_i(\mathcal{X}_i, r_i) = (V_i, E_i)$ be a PRM graph embedded in $\mathcal{C}_i^f$, for some point set $\mathcal{X}_i$ and radius $r_i > 0$ (as defined in Section II-B). The TR, denoted by $\widehat{G}(\vec{\mathcal{X}}, \vec{r}) = (\widehat{V}, \widehat{E})$, is the tensor product of $G_1, \dots, G_R$. In particular, each vertex of $\widehat{G}(\vec{\mathcal{X}}, \vec{r})$ describes a simultaneous placement of the R robots, and similarly an edge of $\widehat{G}(\vec{\mathcal{X}}, \vec{r})$ describes a simultaneous motion of the robots. Formally,

- (i) $\widehat{V} = \{(v_1, \dots, v_R) : \forall i, v_i \in V_i\}$, and
- (ii) for two vertices $W = (w_1, \dots, w_R), U = (u_1, \dots, u_R) \in \widehat{V}$, the edge set $\widehat{E}$ contains the edge (W, U) if for all i , $w_i = u_i$ or $(w_i, u_i) \in E_i$.

Note that robots are allowed to stay put, which differs from prevalent definitions of the tensor product of graphs [37–39]. Notice further that by the definition of G_i , the motion described by each edge in $\widehat{E}$ represents a trajectory for the R robots in which the robot-obstacle collisions are avoided. Next we consider a subgraph of the TR in which also robot-robot collisions are avoided. Given an MRMP problem, $\mathbb{M} = (\mathbb{C}^f, \vec{x}^s, \vec{x}^g)$ we will denote as $\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}(\vec{x}^s, \vec{x}^g)$ its TR constructed from the PRM graphs of the individual robots' $G_{\mathcal{M}_i(\mathcal{X}_i, r_i)}$, but where we remove all edges of the TR representing transitions of the robots that are not collision free. We denote by $d(\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}, \vec{x}^s, \vec{x}^g)$ the minimal cost of collision-free trajectories from $\vec{x}^s$ to $\vec{x}^g$ in the graph $\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}$.

C. Multi-robot clearance and completeness

In preparation for defining the multi-robot equivalent of (ε, δ) -completeness, we first define clearance for the multi-robot case. Given a trajectory Σ , recall that we define for each robot its forbidden space at time $\tau \in [0, 1]$ to be its obstacle space $\mathcal{C}_i^o$ and the configurations that will lead to collisions with other robots. Formally, $\mathcal{C}_i^o(\tau) = \mathcal{C}_i^o \cup_{j \neq i} I_i^j(\sigma_j(\tau))$ is the forbidden space for robot i at time τ . Notice that we define $\mathcal{C}_i^o(\tau)$ only for a given trajectory Σ as it depends on the locations of the other robots.

Definition 3 ($\vec{\delta}$ -clearance). Given a trajectory Σ , we say that Σ has $\vec{\delta}$ -clearance for $\vec{\delta} = (\delta_1, \dots, \delta_R)$ if for each robot i , $1 \leq i \leq R$ and at any time $\tau \in [0, 1]$, the distance from i to the obstacles and to each robot $j \neq i$ is at least δ_i . Formally, Σ has $\vec{\delta}$ -clearance if, for all $1 \leq i \leq R$, $0 \leq \tau \leq 1$, it holds that $\|\sigma_i(\tau) - x\| > \delta_i$, for every $x \in \mathcal{C}_i^o(\tau)$.

Next we define the equivalent of single-robot (ε, δ) -completeness for MRMP:

Definition 4 (Multi-robot $(\varepsilon, \vec{\delta})$ -completeness). Given R robots, a stretch parameter $\varepsilon > 0$, a vector of R sample sets $\vec{\mathcal{X}} = (\mathcal{X}_1, \dots, \mathcal{X}_R)$, and a vector of R connection radii $\vec{r} = (r_1, \dots, r_R)$, we say that the pair $(\vec{\mathcal{X}}, \vec{r})$ is $(\varepsilon, \vec{\delta})$ -complete if for every $\vec{\delta}$ -clear $\mathbb{M} = (\mathbb{C}^f, \vec{x}^s, \vec{x}^g)$ it holds that

$$d(\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}, \vec{x}^s, \vec{x}^g) \leq (1 + \varepsilon) \text{OPT}_{\vec{\delta}},$$

where $\text{OPT}_{\vec{\delta}}$ is the minimal cost of a $\vec{\delta}$ -clear $\mathbb{M}$ solution.

We are ready to state our main contribution.

Theorem 2 (Sufficient conditions for MRMP $(\varepsilon, \vec{\delta})$ -completeness). Let $\varepsilon > 0$ be a stretch factor, let $\vec{\delta}$ be a clearance vector $(\delta_1, \dots, \delta_R)$, and denote $\omega = \frac{\varepsilon}{2(\varepsilon+2)}$. Define the sampling distributions $\vec{\mathcal{X}} = (\mathcal{X}_1, \dots, \mathcal{X}_R)$ and radii vector $\vec{r} = (r_1, \dots, r_R)$, as

$$\mathcal{X}_i = \mathcal{X}_{\omega\delta_i, \delta_i}, \quad r_i = \delta_i(\varepsilon + 1)/(\varepsilon + 2),$$

for every robot $1 \leq i \leq R$. Then $(\vec{\mathcal{X}}, \vec{r})$ is $(\varepsilon, \vec{\delta})$ -complete.

We provide a sketch of the proof. The full proof can be found in the extended version of the paper [34].

Sketch of proof. Let $\Sigma^* = (\sigma_1^*, \dots, \sigma_R^*)$ be a $\vec{\delta}$ -clear solution such that $\text{cost}(\Sigma^*) = \text{OPT}_{\vec{\delta}}$. For every $1 \leq i \leq R$, define $\mathcal{X}_i = \mathcal{X}_{\beta_i, \delta_i}$, and set $\beta_i = \omega\delta_i$. Let $G_i = (V_i, E_i)$ be the PRM graph for robot i using $(\mathcal{X}_i, r_i)$, that is $G_i = G_{\mathcal{M}_i(\mathcal{X}_i, r_i)}$ for $\mathcal{M}_i = (\mathcal{C}_i, x_i^s, x_i^g)$. Let $\widehat{G}$ be the tensor roadmap of $G_1, \dots, G_R$.

In the extended version [34, Lemma 3], we show that for every robot i , G_i contains a collision-free path $\bar{\sigma}_i$ (with respect to the obstacles) from x_i^s to x_i^g , such that $\text{cost}(\bar{\sigma}_i) \leq (1 + \varepsilon)\|\sigma_i^*\|$, and hence $\text{cost}(\bar{\Sigma}) \leq (1 + \varepsilon)\text{OPT}_{\vec{\delta}}$, for $\bar{\Sigma} = (\bar{\sigma}_1, \dots, \bar{\sigma}_R)$. Although each trajectory $\bar{\sigma}_i$ is collision free with respect to $\mathcal{C}_i^f$, it is not necessarily true that $\bar{\Sigma}$ avoids robot-robot collisions. Nevertheless, we show that we can adjust the robots' timing along the trajectories $\bar{\sigma}_1, \dots, \bar{\sigma}_R$, to induce a collision-free trajectory over $\widehat{G}$, which we denote by $\widehat{\Sigma} = (\widehat{\sigma}_1, \dots, \widehat{\sigma}_R)$. As our cost metric is total path length, this timing adjustment does not increase the solution cost.

To do so, we use details from the proof of Theorem 1 (provided in the extended version [34]). In particular, for every robot i , we can pick the trajectory $\bar{\sigma}_i$, which we used above, such that if $Z_i = (z_0^i, \dots, z_{\ell_i}^i)$ denotes the sequence of vertices of G_i visited along $\bar{\sigma}_i$, and $(\tau_0^i, \dots, \tau_{\ell_i}^i)$ denote the time of visitation of those vertices, i.e., $z_j^i = \bar{\sigma}_i(\tau_j^i)$, then the following properties hold for $\rho_i = \delta_i/(\varepsilon + 2)$:

- (i') $\|\sigma_i^*(\tau_k^i) - \sigma_i^*(\tau_{k-1}^i)\| \leq \rho_i, \forall 1 \leq k \leq \ell_i$;
- (ii') $\|z_k^i - \sigma_i^*(\tau_k^i)\| \leq \beta_i, \forall 0 \leq k \leq \ell_i$;
- (iii') $\|z - \sigma_i^*(\tau_k^i)\| \leq \beta_i + \rho_i \leq \delta_i/2, \forall 1 \leq k \leq \ell_i, z \in \text{CH}(z_{k-1}^i, z_k^i)$.

To exploit those properties we define a list $\mathcal{L}$ of triplets of the following form: $\mathcal{L} := \bigcup_{i=1}^R \bigcup_{j=1}^{\ell_i} \{(i, \tau_j^i, z_j^i)\}$. That is, $\mathcal{L}$ contains for every robot $1 \leq i \leq R$, ℓ_i triplets of the form (i, τ_j^i, z_j^i) , where $\tau_j^i \in T_i$ is a timestamp, and $z_j^i \in Z_i$ is the corresponding configuration. Additionally, define $\mathcal{L}^o$ to be a permutation of $\mathcal{L}$, where the triplets are ordered according to the timestamp. That is, $\mathcal{L}^o := \{(i_1, \tau_{i_1}, z_{i_1}), \dots, (i_\ell, \tau_{i_\ell}, z_{i_\ell})\}$, where $\ell = \sum_{i=1}^R \ell_i$, $(i_j, \tau_{i_j}, z_{i_j}) \in \mathcal{L}$ for every $1 \leq j \leq \ell$, and $\tau_{i_j} \leq \tau_{i_{j+1}}$ for every $1 \leq j \leq \ell - 1$.

We describe an iterative scheme that uses $\mathcal{L}^o$ for generating a sequence of composite vertices $V_0, V_1, \dots, V_\ell \in \widehat{V}$, such that $(V_j, V_{j+1}) \in \widehat{E}$. First, define $V_0 = \vec{x}^s$. Next, given that $V_j = (v_{j1}, \dots, v_{jR})$ has already been defined for some $1 \leq j \leq \ell - 1$, set $V_{j+1} = (v_{(j+1)1}, \dots, v_{(j+1)R})$, where $v_{(j+1)i_{j+1}} := z_{i_{j+1}}$, and $v_{(j+1)i'} := v_{ji'}$ for every $i' \neq i_{j+1}$. Namely, when transitioning from V_j to V_{j+1} all the robots stay put, besides robot i_{j+1} whose timestamp appeared in item $j + 1$ of $\mathcal{L}^o$. (An example for $\mathcal{L}$ and $\mathcal{L}^o$ can be found in the extended version [34].)

To complete the proof, we first note that $(V_j, V_{j+1}) \in \widehat{E}$ for every $1 \leq j \leq \ell - 1$. This follows from the values of connection radii $\vec{r}$ we picked, and the fact that $\bar{\sigma}_i$ is obstacle collision-free, guaranteeing that $\{v_j^i, v_{j+1}^i\} \in E_i$ for every robot i , where $(v_1, \dots, v_R) := V_j, (v'_1, \dots, v'_R) := V_{j+1}$.

It remains to prove that the robots do not collide with one another while they move along the path represented by any such edge (V_j, V_{j+1}) . First, recall that there is exactly one robot moving for (V_j, V_{j+1}) . In particular, this is the robot whose index is i_j , which is the first value of the j th triplet $(i_j, \tau_{i_j}, z_{i_j}) \in \mathcal{L}^o$. That is $v_{i_j} \neq v'_{i'_j}$, whereas for any other robot $k \neq i_j$ it holds that $v_k = v'_k$. Note that given two stationary robots $k_1, k_2 \neq i_j$, and assuming that they did

not collide along the previous edge (V_{j-1}, V_j) , they will not collide with each other for (V_j, V_{j+1}) as well.

Next, we prove that robot i_j does not collide with any stationary robot $k \neq i_j$, while moving from v_{i_j} to v'_{i_j} . By definition of $V_0, \dots, V_\ell$, robot k is located at $v_k = z_{j'}^k \in Z_k$ such that $\tau_{j'}^k \leq \tau_{i_j} \leq \tau_{j'+1}^k$, for some $0 \leq j' \leq \ell_k - 1$. Next, recall that $\sigma_k^*, \sigma_{i_j}^*$ are δ_k - and δ_{i_j} -clear, respectively. Thus, it suffices to prove that $\|\sigma_{i_j}^*(\tau_{i_j}) - p\| \leq \delta_{i_j}/2$ for every $p \in \text{CH}(v_{i_j}, v'_{i_j})$ and $\|\sigma_k^*(\tau_{i_j}) - z_{j'}^k\| \leq \delta_k/2$.

Indeed, Property (iii') implies that for $p \in \text{CH}(v_{i_j}, v'_{i_j})$ we have that $\|p - \sigma_{i_j}^*(\tau_{i_j})\| \leq \beta_{i_j} + \rho_{i_j} \leq \delta_{i_j}/2$. Using the triangle inequality and Properties (i') and (ii') we have that $\|z_{j'}^k - \sigma_k^*(\tau_{i_j})\| \leq \|z_{j'}^k - \sigma_k^*(\tau_{j'}^k)\| + \|\sigma_k^*(\tau_{j'}^k) - \sigma_k^*(\tau_{i_j})\| \leq \beta_k + \rho_k \leq \delta_k/2$,

which concludes the proof sketch. $\square$

We emphasize that even though our proof finds a trajectory which uses edges where a single robot moves at a time, the solution that would be found in practice is not necessarily restricted to individual-robot moves. This is due to the fact the tensor roadmap also includes edges representing simultaneous motion of several robots.

D. Discussion

Theorem 2 implies that if a given MRMP planner is guaranteed to find an optimal collision-free path over a TR, then it is also guaranteed to find a $(1 + \varepsilon)$ -approximation of the optimal δ -clear trajectory (in the continuous domain), when each PRM graph is constructed using the sample set $\mathcal{X}_{\omega\delta_i, \delta_i}$ and radius $r_i = \delta_i(\varepsilon + 1)/(\varepsilon + 2)$ for each robot $1 \leq i \leq R$, where $\omega = \varepsilon/(2(\varepsilon + 2))$.

This statement applies, for instance, to M^* and MC-CBS. The former can be viewed as a refined version of A^* for searching the TR. The latter implicitly explores the TR by incrementally considering combinations of single-robot trajectories induced by the PRM graphs, until a combination that yields a collision-free composite trajectory is found. The dRRT* planner implicitly explores the TR via an RRT-style random exploration using a secondary sampling procedure which is employed after the PRM graphs are constructed. Due to this additional randomization step, dRRT* achieves a $(1 + \varepsilon)$ -approximation only *asymptotically*. Nevertheless, our analysis simplifies the usage of this algorithm by derandomizing the construction of PRM graphs used in dRRT*.

Finally, we provide an example for the number of samples that should be used according to Theorem 2 within each PRM roadmap for specific parameters. In particular, we report in the table below the value $|\mathcal{X}_{\omega\delta_i, \delta_i}|$, for varying values of the stretch parameter ε , dimension d , and clearance vector $(\delta_1, \dots, \delta_R)$, where $\delta_i = 0.1$ for all robots. While these values are quite large, particularly for higher dimensions, we emphasize that we do not expect our bounds to be tight, as observed in Section IV. This suggests that smaller sample sets are sufficient for $(\varepsilon, \vec{\delta})$ -completeness, which we aim to explore in future research (see Section V).

IV. EXPERIMENTAL RESULTS

We provide experimental results to support our theoretical findings, focusing on the case of multiple disc robots

d	$\varepsilon = \infty$	$\varepsilon = 5$	$\varepsilon = 1$	$\varepsilon = 0.5$	$\varepsilon = 0.25$
2	181	313	1201	3281	$1.05 \cdot 10^4$
3	2331	6119	$5.68 \cdot 10^4$	$2.43 \cdot 10^5$	$1.43 \cdot 10^6$
4	$4.93 \cdot 10^4$	$1.49 \cdot 10^5$	$2.83 \cdot 10^6$	$2.19 \cdot 10^7$	$2.21 \cdot 10^8$
5	$9.09 \cdot 10^5$	$4.37 \cdot 10^6$	$1.69 \cdot 10^8$	$2.23 \cdot 10^9$	$3.94 \cdot 10^{10}$
6	$1.89 \cdot 10^7$	$1.5 \cdot 10^8$	$1.18 \cdot 10^{10}$	$2.46 \cdot 10^{11}$	$7.82 \cdot 10^{12}$

operating in a planar domain. We study the effect that the stretch parameter ε , which determines the structure of the underlying PRM graphs within the tensor product graph $\hat{G}$ in Theorem 2, has on the actual solution quality. We observe that the resulting approximation factor is in fact significantly lower than $1 + \varepsilon$. This suggests that our analysis can be further refined to support even sampling distributions with fewer samples (see Section V).

Concerning the design of the experiments, we were faced with two challenges. First, current MRMP algorithms [6, 11, 32] (including our own) are limited in their ability to cope with the large tensor roadmaps required to guarantee path quality according to our analysis; this curbed our ability to go to large numbers of robots in the experiments, and strongly motivates further improving of such algorithms. Secondly, in order to provide meaningful experimental reports, we need yardsticks to compare to; the problem is that optimal MRMP algorithms are not known to be tractable even for the simple case of two unit disc robots moving amid obstacles in the plane, and it is highly non-trivial to calculate optimal solutions when coordination is required. We explain below how we overcame this latter impediment.

A. Scenarios

The scenarios are illustrated in Figure 2. The first scenario, [Left], consists of two robots in an obstacle-free environment. We use this simple example to benchmark our solution against an optimal δ -clear solution using a recent work that provides characterization of optimal trajectories for two disc robots *in the absence of obstacles* [40]. The [Center-left] scenario uses a more complicated workspace topology, which potentially requires more samples to achieve a near-optimal solution, since multiple straight-line segments are required to approximate every single-robot trajectory. The third scenario, [Center-right], which consists of four robots, aims to test our theory for a tight setting, which is obtained by tightly packing the four robots inside a circular barrier. The fourth scenario, [Right], which consists of seven robots, aims to test our theory for a larger number of robots, where additional coordination is required to achieve a solution.

B. Results

To test our theory on the aforementioned scenarios, we constructed PRM graphs using our staggered grid as the sample set, corresponding to different values of the stretch parameter ε and clearance δ (for simplicity, we use the same value of δ across all robots, i.e., $\vec{\delta} := \{\delta, \dots, \delta\}$). For all the scenarios, we set the value δ to be equal to the *static clearance* μ (see caption of Figure 2). We set the stretch parameter ε to different values in the range $[0.75, \infty)$. We then use A^* -search for the first three scenarios, and MC-CBS for the seven-robot scenario, to obtain the best solution from the resulting tensor roadmap.

The results are reported in Figure 3, where we plot the approximation ratio obtained using the staggered grid $\mathcal{X}_{\omega\delta, \delta}$ set to guarantee an approximation factor of at most $1 + \varepsilon$ (see




Fig. 2: Test scenarios for multiple disc robots, where a circle and a disc of the same color represent the start and goal positions, respectively, of a robot. [Left] A 2-robot obstacle-free scenario. The optimal trajectories, which were derived in [40], are drawn for each robot. The robots’ radius is 0.09 and the static clearance μ is equal to 0.02. The latter value measures the minimum over (i) the inter-robot distances at the initial placement, (ii) the distance of each robot at the initial placement from the obstacles, and (iii,iv) the respective quantities for the target placement. [Center-left] A 2-robot scenario with a spiral obstacle with robots’ radius 0.06 and static clearance $\mu = 0.04$. [Center-right] Four robots tightly placed within a circular barrier, using robot radius 0.19 and static clearance $\mu = 0.02$. [Right] A 7-robot scenario with robots’ radius 0.08 and the static clearance $\mu = 0.04$.




Fig. 3: We report for each of the four scenarios its approximation factors, which are represented by the ratio between the cost of the solution obtained from the tensor roadmap and OPT_δ . Notice that the tensor-roadmap solution is not necessarily δ -clear and thus can be of lower cost than OPT_δ , which explains why two of the plots get approximation factors smaller than 1. Next to the points in each plot is the number of collision-free samples in the PRM graph of an individual robot. Due to the prohibitive running times incurred by MC-CBS on the 7-robot scenario and by our Python implementations of A* on the 4-robot scenario, we report the solution quality for these scenarios only for a subset of values of ϵ .

Theorem 2). The reported approximation factor represents the ratio between the cost of the best solution obtained from the tensor roadmap, denoted by $\hat{\Sigma}$, and the optimal δ -clear solution whose cost is denoted by OPT_δ . For the first scenario we obtain the value OPT_δ using [40], as we reported earlier. For the [Center-left] scenario, OPT_δ is equal to the sum of the shortest δ -clear trajectories for the two individual robots, as one of the robots can move after the other finishes its motion without increasing the overall cost of the solution. For the [Center-right] scenario, OPT_δ is equal to the perimeter of the circle going through the robot centers at the initial positions (as each robot traverses a quarter of the circle). For the [Right] scenario, OPT_δ is equal to the sum of the Euclidean distances between each robot’s origin and destination, as the robots can move one after the other to obtain an optimal solution.

In all the experiments we obtain an approximation factor that is significantly lower than what our worst-case analysis predicts. For instance, already when setting the stretch parameter to $\epsilon = 50$ we obtain an approximation factor of at

most 1.075. Moreover, for $\epsilon \leq 1.5$ we obtain approximation factors below 1 in both two-robot scenarios. This is possible as the solution obtained from the tensor roadmap is not necessarily δ -clear, which allows robots to take shortcuts in proximity to obstacles and each other. The overall trend of the graphs complies with our expectation: when the stretch parameter decreases, we obtain improved solutions. The biggest improvement (at least for the two-robot scenarios) occurs when ϵ goes below 2, due to the rapid increase in the number of samples in $\mathcal{X}_{\omega, \delta, \epsilon}$. In some cases using a smaller number of samples may yield better solutions, as in the 2-robot spiral scenario for stretch factors $\epsilon = 50$ and $\epsilon = 20$. The explanation is that the smaller sample set gets closer to the (approximate) optimal solution by chance. Still, the *worst-case approximation factor* is guaranteed to improve as the size of the staggered grid increases.

V. DISCUSSION AND FUTURE WORK

We developed sufficient theoretical conditions for finite-sample near-optimality of the tensor roadmap, which is an underlying structure in several sampling-based algorithms for MRMP. We also presented a new sampling-scheme, termed the staggered grid, for near-optimal motion planning for individual robots, which requires fewer samples than previous work.

Our work raises interesting questions for further investigation both in practice and theory. The scalability issues we encountered when testing our theoretical finding on MC-CBS—a continuous extension of a state-of-the-art method for MAPF—motivate the study of more effective methods for exploring tensor roadmaps introduced by large PRM graphs. On the positive side, the scenarios that we did manage to solve suggest that near-optimality can be achieved with smaller sample sets than our theory prescribes. This motivates the development of even more compact sampling-distributions for the single-robot case, and refining our proof technique for the multi-robot case (Theorem 2).

ACKNOWLEDGMENTS

The authors thank Matt Tsao for fruitful discussions, Nir Goren for providing a Python motion-planning framework implementation, and Jiaoyang Li for providing an implementation for MC-CBS.

REFERENCES

- [1] K. Solovey and D. Halperin, "On the hardness of unlabeled multi-robot motion planning," *International Journal of Robotic Research*, vol. 35, no. 14, pp. 1750–1759, 2016.
- [2] P. G. Spirakis and C.-K. Yap, "Strong NP-hardness of moving many discs," *Information Processing Letters*, vol. 19, no. 1, pp. 55–59, 1984.
- [3] J. E. Hopcroft, J. T. Schwartz, and M. Sharir, "On the complexity of motion planning for multiple independent objects; PSPACE-hardness of the 'Warehouseman's problem'," *International Journal of Robotics Research*, vol. 3, no. 4, pp. 76–88, 1984.
- [4] J. K. Johnson, "On the relationship between dynamics and complexity in multi-agent collision avoidance," *Autonomous Robots*, vol. 42, no. 7, pp. 1389–1404, 2018.
- [5] S. D. Han and J. Yu, "Effective heuristics for multi-robot path planning in warehouse environments," in *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. IEEE, 2019, pp. 10–12.
- [6] G. Wagner and H. Choset, "Subdimensional expansion for multirobot path planning," *Artificial Intelligence*, vol. 219, pp. 1–24, 2015.
- [7] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [8] H. Ma, W. Hönig, T. K. S. Kumar, N. Ayanian, and S. Koenig, "Lifelong path planning with kinematic constraints for multi-agent pickup and delivery," in *Conference on Artificial Intelligence, AAAI*, 2019, pp. 7651–7658.
- [9] S. Choudhury, K. Solovey, M. J. Kochenderfer, and M. Pavone, "Efficient large-scale multi-drone delivery using transit networks," in *International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4543–4550.
- [10] W. Hönig, T. K. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig, "Summary: Multi-agent path finding with kinematic constraints," in *International Joint Conference on Artificial Intelligence*, 2017, pp. 4869–4873.
- [11] J. Li, P. Surynek, A. Felner, H. Ma, T. K. S. Kumar, and S. Koenig, "Multi-agent path finding for large agents," in *Conference on Artificial Intelligence AAAI*, 2019, pp. 7627–7634.
- [12] I. Solis, R. Sandstrom, J. Motes, and N. M. Amato, "Roadmap-optimal multi-robot motion planning using conflict-based search," *CoRR*, vol. abs/1909.13352, 2019.
- [13] A. Adler, M. de Berg, D. Halperin, and K. Solovey, "Efficient multi-robot motion planning for unlabeled discs in simple polygons," *IEEE Trans. Automation Science and Engineering*, vol. 12, no. 4, pp. 1309–1317, 2015.
- [14] I. Solomon and D. Halperin, "Motion planning for multiple unit-ball robots in $\mathbb{R}^d$," in *Workshop on the Algorithmic Foundations of Robotics, WAFR*, 2018, pp. 799–816.
- [15] M. Turpin, N. Michael, and V. Kumar, "Trajectory planning and assignment in multirobot systems," in *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2012, pp. 175–190.
- [16] K. Solovey, J. Yu, O. Zamir, and D. Halperin, "Motion planning for unlabeled discs with optimality guarantees," in *Robotics: Science and Systems*, 2015.
- [17] L. E. Kavraki, P. Švestka, J. . Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [18] L. E. Kavraki, M. N. Kolountzakis, and J. . Latombe, "Analysis of probabilistic roadmaps for path planning," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 1, pp. 166–171, 1998.
- [19] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [20] J. J. Kuffner and S. M. LaValle, "RRT-Connect: An efficient approach to single-query path planning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2000, pp. 995–1001.
- [21] M. Kleinbort, K. Solovey, Z. Littlefield, K. E. Bekris, and D. Halperin, "Probabilistic completeness of RRT for geometric and kinodynamic planning with forward propagation," *IEEE Robotics Autom. Lett.*, vol. 4, no. 2, pp. 277–283, 2019. [Online]. Available: <https://doi.org/10.1109/LRA.2018.2888947>
- [22] L. Janson, E. Schmerling, A. Clark, and M. Pavone, "Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 883–921, 2015.
- [23] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [24] K. Solovey, L. Janson, E. Schmerling, E. Frazzoli, and M. Pavone, "Revisiting the asymptotic optimality of RRT*," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 2189–2195.
- [25] K. Hauser and Y. Zhou, "Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space," *IEEE Trans. Robotics*, vol. 32, no. 6, pp. 1431–1443, 2016.
- [26] M. Kleinbort, E. Granados, K. Solovey, R. Bonalli, K. E. Bekris, and D. Halperin, "Refined analysis of asymptotically-optimal kinodynamic planning in the state-cost space," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 6344–6350.
- [27] G. Sánchez-Ante and J. Latombe, "Using a PRM planner to compare centralized and decoupled planning for multi-robot systems," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2002, pp. 2112–2119.
- [28] P. Švestka and M. H. Overmars, "Coordinated path planning for multiple robots," *Robotics and Autonomous Systems*, vol. 23, no. 3, pp. 125–152, 1998.
- [29] K. Solovey and D. Halperin, " k -Color multi-robot motion planning," *International Journal of Robotic Research*, vol. 33, no. 1, pp. 82–97, 2014.
- [30] A. Krontiris, R. Shome, A. Dobson, A. Kimmel, and K. E. Bekris, "Rearranging similar objects with a manipulator using pebble graphs," in *IEEE-RAS International Conference on Humanoid Robots, Humanoids*, 2014, pp. 1081–1087.
- [31] K. Solovey, O. Salzman, and D. Halperin, "Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 501–513, 2016.
- [32] R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris, "dRRT*: Scalable and informed asymptotically-optimal multi-robot motion planning," *Auton. Robots*, vol. 44, no. 3–4, pp. 443–467, 2020.
- [33] M. Tsao, K. Solovey, and M. Pavone, "Sample complexity of probabilistic roadmaps via epsilon-nets," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [34] D. Dayan, K. Solovey, M. Pavone, and D. Halperin, "Near-optimal multi-robot motion planning with finite sampling," *CoRR*, vol. abs/2011.08944, 2020. [Online]. Available: <https://arxiv.org/abs/2011.08944>
- [35] K. Solovey and M. Kleinbort, "The critical radius in sampling-based motion planning," *International Journal of Robotics Research*, vol. 39, no. 2–3, pp. 266–285, 2020.
- [36] D. Halperin and M. Sharir, "Arrangements," in *Handbook of Discrete and Computational Geometry*, 3rd ed., J. E. Goodman, J. O'Rourke, and C. Tóth, Eds. Chapman & Hall/CRC, 2017, ch. 28, pp. 1343–1376.
- [37] R. Hammack, W. Imrich, and S. Klavzar, *Handbook of Product Graphs, Second Edition*, 2nd ed. USA: CRC Press, Inc., 2011.
- [38] R. Balakrishnan and P. Paulraja, "Hamilton cycles in tensor product of graphs," *Discrete Mathematics*, vol. 186, no. 1, pp. 1 – 13, 1998.
- [39] Shitov, "Counterexamples to Hedetniemi's conjecture," *Annals of Mathematics*, vol. 190, no. 2, p. 663, 2019.
- [40] D. G. Kirkpatrick and P. Liu, "Characterizing minimum-length coordinated motions for two discs," *CoRR*, vol. abs/1607.04005, 2016.